



Microsoft

Exam Questions 70-464

Developing Microsoft SQL Server 2012 Databases

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 build a stored procedure that amortizes the invoice amount. Which code segment should you use to create the stored procedure? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

| Ordered List Title | Answer Choices Title |
|--|---|
| | <pre> RECONFIGURE; EXEC sp_configure 'clr enabled', '1'; EXEC sp_recompleie @objname = 'TaxCalc' CREATE PROCEDURE Accounting.Amortize(@total decimal(8,2), @period int) RETURNS decimal(8,2) AS EXTERNAL NAME TaxCalc.TreyResearch.Amortize; CREATE ASSEMBLY TaxCalc FROM 'C:\temp\TreyTax.DLL' CREATE ASSEMBLY TaxCalc FROM 'C:\temp\Amortize.cs' </pre> |
| <p><< Move Remove >></p> | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

D:\Documents and Settings\useralbo\Desktop\1.jpg
<http://msdn.microsoft.com/en-us/library/ms131089.aspx> <http://msdn.microsoft.com/en-us/library/ms131048.aspx> <http://msdn.microsoft.com/en-us/library/ms187926.aspx>

NEW QUESTION 4

- (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 5

- (Exam Topic 2)

While testing usp_GetFutureSessions, you discover that IX_Sessions is accessed by a scan rather than a seek. You need to minimize the amount of time it takes to execute usp_GetFutureSessions. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Change line 02 of Indexes.sql to:
`(DeliveryTime, SessionID)`
- B. At line 04 of Indexes.sql, add:
`WHERE GETDATE() < DeliveryTime;`
- C. Change line 02 of Indexes.sql to:
`(SpeakerID, RoomID, DeliveryTime)`
- D. Change line 74 of Procedures.sql to:
`WHERE GETDATE() > DeliveryTime;`
- E. Change line 74 of Procedures.sql to:
`WHERE GETDATE() < DeliveryTime;`
- F. At line 04 of Indexes.sql, add:
`WHERE GETDATE() > DeliveryTime;`

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

Answer: BE

Explanation:

Future delivery dates.

NEW QUESTION 6

- (Exam Topic 2)

You need to modify usp_SelectSpeakersByName to support server-side paging. The solution must minimize the amount of development effort required. What should you add to usp_SelectSpeakersByName?

- A. A table variable
- B. An OFFSET-FETCH clause
- C. The ROWNUMBER keyword
- D. A recursive common table expression

Answer: B

Explanation:

<http://www.mssqltips.com/sqlservertip/2696/comparing-performance-for-different-sql-server-paging-methods/> <http://msdn.microsoft.com/en-us/library/ms188385.aspx>

<http://msdn.microsoft.com/en-us/library/ms180152.aspx> <http://msdn.microsoft.com/en-us/library/ms186243.aspx> <http://msdn.microsoft.com/en-us/library/ms186734.aspx>

<http://www.sqlserver-training.com/how-to-use-offset-fetch-option-in-sql-server-order-by-clause/>

http://www.sqlservercentral.com/blogs/juggling_with_sql/2011/11/30/using-offset-and-fetch/

NEW QUESTION 7

- (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 8

- (Exam Topic 2)

You are evaluating the table design.

You need to recommend a change to Tables.sql that reduces the amount of time it takes for usp_AttendeesReport to execute.

What should you add at line 14 of Tables.sql?

- A. FullName nvarchar(100) NOT NULL CONSTRAINT DF_FullName DEFAULT (dbo.CreateFullName (FirstName, LastName)),
- B. FullName AS (FirstName + ' ' + LastName),
- C. FullName nvarchar(100) NOT NULL DEFAULT (dbo.CreateFullName (FirstName, LastName)).
- D. FullName AS (FirstName + ' ' + LastName) PERSISTED,

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms188300.aspx> <http://msdn.microsoft.com/en-us/library/ms191250.aspx>

NEW QUESTION 9

- (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 10

- (Exam Topic 2)

You need to add a new column named Confirmed to the Attendees table. The solution must meet the following requirements:

- Have a default value of false.
- Minimize the amount of disk space used.

Which code block should you use?

- A. ALTER TABLE Attendees ADD Confirmed bit DEFAULT 0;
- B. ALTER TABLE Attendees ADD Confirmed char(1) DEFAULT '1';
- C. ALTER TABLE Attendees ADD Confirmed bit DEFAULT 1;
- D. ALTER TABLE Attendees ADD Confirmed char(1) DEFAULT '1';

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms177603.aspx>

NEW QUESTION 10

- (Exam Topic 2)

You are evaluating the index design.

You need to recommend a change to Indexes.sql that will minimize the amount of time it takes for usp_AttendeesReport to execute. The solution must minimize the amount of database fragmentation.

Which line of code should you use to replace line 12 of Indexes.sql?

- A. (LastName);
- B. (FirstName) INCLUDE (LastName);
- C. (LastName, FirstName);
- D. (LastName) INCLUDE (FirstName);

Answer: C

NEW QUESTION 15

- (Exam Topic 3)

You need to recommend a solution to ensure that SQL1 supports the auditing requirements of usp_UpdateEmployeeName. What should you include in the recommendation?

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

Answer: D

NEW QUESTION 17

- (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 20

- (Exam Topic 4)

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

- A. OPEN PRODUCTSCERT;
ALTER PROCEDURE Production.ProductDetails_Insert
WITH ENCRYPTION;
CLOSE PRODUCTSCERT;
- B. OPEN DBCERT;
ALTER PROCEDURE Production.ProductDetails_Insert
WITH ENCRYPTION;
CLOSE DBCERT;
- C. ADD SIGNATURE TO Production.ProductDetails_Insert
BY CERTIFICATE DBCERT;
- D. ADD SIGNATURE TO Production.ProductDetails_Insert
BY CERTIFICATE PRODUCTSCERT;

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

D. Option D

Answer: C

Explanation:

<http://msdn.microsoft.com/en-us/library/bb669102.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 27

- (Exam Topic 4)

An administrator provides a digital certificate named ServerCert.

You need to implement Transparent Data Encryption (TDE) on ProductsDB. Which code segment should you use?

- A. USE PRODUCTSDB; GO CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = TRIPLE_DES_3KEY ENCRYPTION BY SERVER CERTIFICATE DBCERT; GO ALTER DATABASE PRODUCTSDB SET ENCRYPTION ON; GO
- B. USE PRODUCTSDB; GO CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = TRIPLE_DES_3KEY ENCRYPTION BY SERVER CERTIFICATE PRODUCTSCERT; GO ALTER DATABASE PRODUCTSDB SET ENCRYPTION ON; GO
- C. USE PRODUCTSDB; GO CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES_256 ENCRYPTION BY SERVER CERTIFICATE PRODUCTSCERT; GO ALTER DATABASE PRODUCTSDB SET ENCRYPTION ON; GO
- D. USE PRODUCTSDB; GO CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES_256 ENCRYPTION BY SERVER CERTIFICATE DBCERT; GO ALTER DATABASE PRODUCTSDB SET ENCRYPTION ON; GO

Answer: C

Explanation:

<http://msdn.microsoft.com/en-us/library/bb934049.aspx>

NEW QUESTION 32

- (Exam Topic 5)

You need to resolve the performance issues of the usp_getOpenings stored procedure. Which three actions should you perform? Each correct answer presents part of the solution.

- A. Delete lines 05 through 08
- B. Replace lines 12, 13, and 14 with the Transact-SQL segment: WHERE (CONTAINS(o.Description, 'ISABOUT(' + @keyword + ' weight (.5)))
- C. Create a full text index on the Description column

- D. Replace lines 12, 13, and 14 with the Transact_SQL segment: WHERE (CONTAINS(o.Description, @keyword))
- E. Replace lines 12, 13, and 14 with the Transact SQL Segment:WHERE (Contains(o.Description, 'FORMSOF(INFLECTIONAL, '+@keyword+')))

Answer: ACE

Explanation:

Scenario: You also discover that the usp_GetOpenings stored procedure takes a long time to run and that the non-clustered index on the Description column is not being used.

The FORMSOF term performs matches using other linguistic forms of the word. The following is the FORMSOF term syntax:

FORMSOF (<generation_type>,<match_words>)

The generation type specifies how Microsoft Windows Search chooses the alternative word forms. The INFLECTIONAL value chooses alternative inflection forms for the match words. If the word is a verb, alternative tenses are used. If the word is a noun, the singular, plural, and possessive forms are used to detect matches.

References:

<https://docs.microsoft.com/en-us/windows/win32/search/-search-sql-formsof>

NEW QUESTION 33

- (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 35

- (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 36

- (Exam Topic 5)

You need to design a solution that meets the refactoring requirements. Which type of object should you include in the solution?

- A. An indexed view
- B. An aggregate function
- C. A distributed view
- D. A table-valued function

Answer: D

NEW QUESTION 37

- (Exam Topic 6)

You need to ensure that usp_AddXMLOrder can be used to validate the XML input from the retailers. Which parameters should you add to usp_AddXMLOrder on line 04 and line 05? (Each correct answer presents part of the solution. Choose all that apply.)

- A. @schema varbinary(100).
- B. @items varchar(max).
- C. @schema sysname.
- D. @items varbinary(max).
- E. @items xml.
- F. @schema xml.

Answer: CE

NEW QUESTION 40

- (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. Add WITH (FORCESEEK) to line 69 in `usp.GetOrdersByProduct`.
- C. Add WITH RECOMPILE to line 64 in `usp.GetOrdersByProduct`.
- D. Execute `sp_recompile usp.GetOrdersByProduct'`.

Answer: B

NEW QUESTION 44

- (Exam Topic 6)

You need to implement a solution that addresses the index monitoring requirements. What should you do?

- A. Schedule a SQL Server Agent job that saves data from the dynamic management views to a table in the database.
- B. Create a SQL Server Audit that saves data to a log file, and then create a SQL Server Audit Specification that gathers data from the DATABASE_OPERATION group.
- C. Create a performance monitor Data Collector Set (DCS) that monitors the SQL Server counters.
- D. Schedule a SQL Server Profiler trace, and then save the trace data to a table in the database.

Answer: A

NEW QUESTION 49

- (Exam Topic 6)

You need to modify `usp_GetOrdersAndItems` to ensure that an order is NOT retrieved by `usp_GetOrdersAndItems` while the order is being updated. What should you add to `usp_GetOrdersAndItems`?

- A. Add SET TRANSACTION ISOLATION LEVEL SERIALIZABLE to line 03.
- B. Add SET TRANSACTION ISOLATION LEVEL SNAPSHOT to line 03.
- C. Add (UPDLOCK) to the end of line 06.
- D. Add (READPAST) to the end of line 06.

Answer: D

NEW QUESTION 50

- (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 52

- (Exam Topic 7)

You need to redesign the system to meet the scalability requirements of the application. 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> , UserId int NOT NULL INDEX ix_UserId NONCLUSTERED HASH WITH (BUCKET_COUNT=2), </pre> | |
| <pre> , UserId int NOT NULL INDEX x_UserId NONCLUSTERED HASH WITH (BUCKET_COUNT=900000), </pre> | |
| <pre> POSLocation int NOT NULL, StatusID int NOT NULL, CreateDate datetime2 NOT NULL, Price money) </pre> | |
| <pre> POSTransactionId int NOT NULL PRIMARY KEY CLUSTERED </pre> | |
| <pre> POSTransactionId int NOT NULL </pre> | |
| <pre> ALTER DATABASE CoffeeTransactions ADD FILEGROUP [CoffeeTransactions_inmem] CONTAINS MEMORY_OPTIMIZED_DATA </pre> | |
| <pre> ON [CoffeeTransactions_inmem] </pre> | |
| <pre> WITH (MEMORY_OPTIMIZED=ON, DURABILITY=SCHEMA_ONLY) </pre> | |
| <pre> POSTransactionId int NOT NULL PRIMARY KEY CLUSTERED HASH WITH (BUCKET_COUNT=1000000) </pre> | |
| <pre> , UserId int NOT NULL NONCLUSTERED INDEX ix_UserId, </pre> | |
| <pre> CREATE TABLE dbo.POSTransaction (</pre> | |
| <pre> POSTransactionId int NOT NULL PRIMARY KEY NONCLUSTERED HASH WITH (BUCKET_COUNT=1) </pre> | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

```

ALTER DATABASE CoffeeTransactions
ADD FILEGROUP [CoffeeTransactions_inmem
] CONTAINS MEMORY_OPTIMIZED_DATA

```

Box 2:

```

CREATE TABLE dbo.POSTransaction (

```

Box 3:

```

,
UserId int NOT NULL
INDEX x_UserId NONCLUSTERED
HASH WITH (BUCKET_COUNT=900000),

```

Box 4:

```

POSTransactionId int NOT NULL
PRIMARY KEY CLUSTERED
HASH WITH (BUCKET_COUNT=1000000)

```

Box 5:

```

POSLocation int NOT NULL,
StatusID int NOT NULL,
CreateDate datetime2 NOT NULL,
Price money

```

Box 6:

```

WITH (MEMORY_OPTIMIZED=ON,
DURABILITY=SCHEMA_ONLY)

```

Box 7:

```

ON [CoffeeTransactions_inmem]

```

Note:

* MEMORY_OPTIMIZED_DATA

First create a memory-optimized data filegroup and add a container to the filegroup. Then create a memory-optimized table.

* You must specify a value for the BUCKET_COUNT parameter when you create the memory-optimized table. In most cases the bucket count should be between 1 and 2 times the number of distinct values in the index key.

* Example:

```

-- create a durable (data will be persisted) memory-optimized table
-- two of the columns are indexed
CREATE TABLE dbo.ShoppingCart (
ShoppingCartId INT IDENTITY(1,1) PRIMARY KEY NONCLUSTERED,
UserId INT NOT NULL INDEX ix_UserId NONCLUSTERED HASH WITH (BUCKET_COUNT=1000000),
CreatedDate DATETIME2 NOT NULL, TotalPrice MONEY
) WITH (MEMORY_OPTIMIZED=ON) GO

```

NEW QUESTION 56

- (Exam Topic 8)

You have a SQL Server 2012 instance.

You plan to create an application that uses spatial data.

You need to create an object that will support the representation of the surface area of all the oceans. Which code segment should you use?

- A. DECLARE @g GEOGRAPHY =
GEOGRAPHY::STGeomFromText (
'FULLGLOBE', 4326
);
- B. DECLARE @g GEOGRAPHY =
GEOGRAPHY::STGeomFromText (
'POLYGON(0 0, 0 10, 10 10, 10 0, 0 0)', 4326
);
- C. DECLARE @g GEOGRAPHY =
GEOGRAPHY::STGeomFromText ('
COMPOUNDCURVE (
CIRCULARSTRING(0 -50, 90 0, 0 50),
CIRCULARSTRING(0 50, 45 50, -90 50),
CIRCULARSTRING(-90 50, 0 0, -90 -50),
CIRCULARSTRING(-90 -50, 45 -50, 0 -50)', 4326
);
- D. DECLARE @g GEOGRAPHY =
GEOGRAPHY::STGeomFromText (
'CIRCULARSTRING(0 50, 90 50, 180 50)', 4326
);

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

Answer: A

NEW QUESTION 61

- (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

Answer Area

```
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
)
```

- 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 62

- (Exam Topic 8)

You have a SQL Server 2012 database named Database1.

Database1 has a table named Customers. Customers contains more than 1 million rows. The database has a stored procedure that was created by using the following script:

```
CREATE PROCEDURE up_customers
    @CustomerTypeList nvarchar(400)
AS
SELECT CustomerID,
    FirstName,
    LastName
FROM dbo.customers
WHERE CustomerTypeID IN (@CustomerTypeList);
```

You need to ensure that up_customers returns rows when the following statement is executed: EXECUTE up_customers'1,2,3,4,5';
 What should you do?

- A. Update @CustomerTypeList to use the int data type.
- B. Convert @CustomerTypeList to a table variable.
- C. Convert @CustomerTypeList to an XML variable.
- D. Update @CustomerTypeList to use the XML data type.

Answer: B

NEW QUESTION 63

- (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 65

- (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 68

- (Exam Topic 8)

Your network contains a server that has SQL Server 2014 installed. You create a table by using the following script:

```
CREATE TABLE dbo.Products
(
    id int NOT NULL,
    ProductName nvarchar(50) NULL,
    ProductManufacturer nvarchar(50) NULL,
    ProductDescription nvarchar(200) NULL,
    CONSTRAINT PK_Products PRIMARY KEY CLUSTERED (id)
)
ON [PRIMARY]
GO
```

You need to recommend a solution to ensure that each combination of ProductName and ProductManufacturer is not duplicated. What should you recommend creating?

- A. A UNIQUE constraint
- B. A filtered index
- C. A columnstore index
- D. A CHECK constraint

Answer: A

NEW QUESTION 71

- (Exam Topic 8)

You have a database hosted on SQL Azure.

You are developing a script to create a view that will be used to update the data in a table.

The following is the relevant portion of the script. (Line numbers are included for reference only.)

```
01 CREATE VIEW View1
02 AS
03 SELECT
04 ...
05 WHERE Column1 = 'City1'
06
```

You need to ensure that the view can update the data in the table, except for the data in Column1. Which code segment should you add at line 06?

- A. WITH CHECK OPTION
- B. WITH VIEW_METADATA
- C. WITH ENCRYPTION
- D. WITH SCHEMABINDING

Answer: A

Explanation:

The question concerning the view that has a clause "WHERE Column1 = 'City1' is wrong. That's not what the CHECK option is made for. Actually you will be able to update ONLY the rows satisfied by that WHERE clause, that is, only the rows with the Column1 being 'City1'.

None of the answers are valid from that question. You need a trigger to achieve that.

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

NEW QUESTION 73

- (Exam Topic 8)

You need to encapsulate a T-SQL script into a reusable user-defined object.

The object must meet the following requirements:

- Permit insertions into a table variable.
- Support structured exception handling.
- Prevent changes to the definition of referenced objects.
- Support the use of the APPLY operator on the output of the object. Which type of object should you use?

- A. An inline table-valued function
- B. A stored procedure
- C. A scalar user-defined function
- D. A multi-statement table-valued function

Answer: C

NEW QUESTION 75

- (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:

- ▶ Identify unused indexes.
- ▶ Identify indexes that need to be defragmented. 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.)

| Functions | Answer Area |
|-----------------------------------|--|
| sys.dm_db_index_usage_stats | Identify unused indexes. |
| sys.dm_db_index_operational_stats | Identify indexes that need to be defragmented. |
| sys.dm_db_index_physical_stats | |
| sys.dm_db_missing_index_columns | |
| sys.dm_db_missing_index_details | |
| sys.dm_db_missing_index_groups | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

* sys.dm_db_index_usage_stats

Returns counts of different types of index operations and the time each type of operation was last performed.

* sys.dm_db_index_physical_stats

Returns size and fragmentation information for the data and indexes of the specified table or view.

NEW QUESTION 80

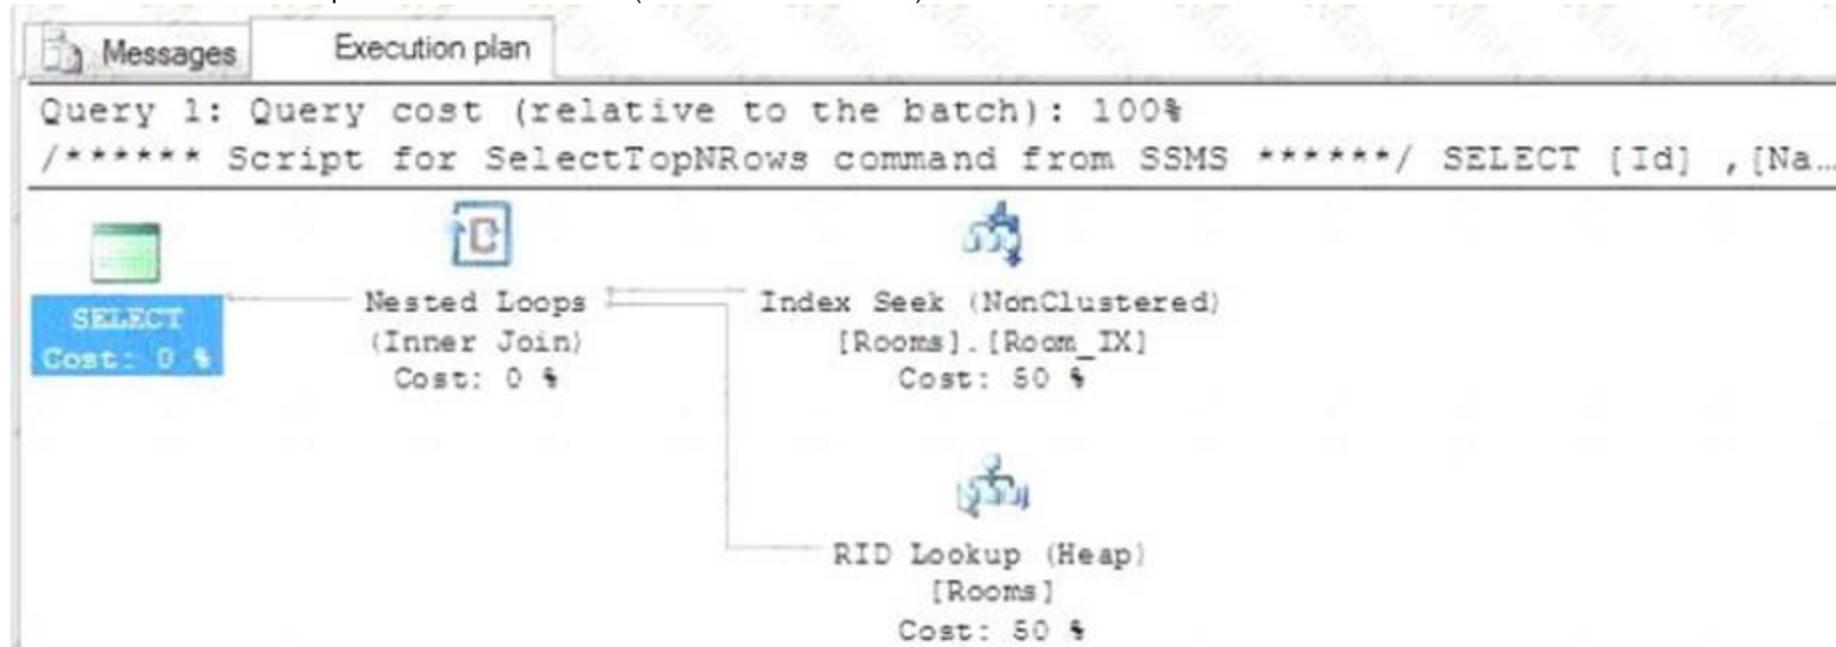
- (Exam Topic 8)

You have a table named Rooms that contains three columns.

You execute the following query:

```
SELECT [Id],
       [RoomName],
       [Position]
FROM [dbo].[Rooms]
WHERE [RoomName] = 'Room1'
```

You discover the execution plan shown in the exhibit. (Click the Exhibit button.)



You need to recommend a solution to reduce the amount of time it takes to execute the query. What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Include the RoomName column and the Position column in the Room_IX index.
- B. Create a nonclustered index for RoomName, Id, and Position.
- C. Create a clustered index for Id.
- D. Use the WITH (INDEX(Room_IX),NOLOCK) query hint.

Answer: B

NEW QUESTION 85

- (Exam Topic 8)

You are designing two stored procedures named Procedure1 and Procedure2. You identify the following requirements:

- ▶ Procedure1 must take a parameter that ensures that multiple rows of data can pass into the stored procedure.
- ▶ Procedure2 must use business logic that resides in a Microsoft .NET Framework assembly. You need to identify the appropriate technology for each stored procedure.

Which technologies should you identify?

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

| Technologies | Answer Area |
|----------------------------------|-------------|
| Common language runtime (CLR) | Procedure1 |
| Extensible Markup Language (XML) | Procedure2 |
| a table-valued parameter (TVP) | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms131102.aspx> <http://msdn.microsoft.com/en-us/library/bb522446.aspx> <http://msdn.microsoft.com/en-us/library/bb510489.aspx>

NEW QUESTION 86

- (Exam Topic 8)

You review a query that runs slowly. The query accesses data in a table named Schema1.Table1. The following is the relevant portion of the execution plan for the query:

```
<MissingIndexes>
  <MissingIndexGroup Impact="95.8296">
    <MissingIndex Database="DB1" Schema="Schema1" Table="Table1">
      <ColumnGroup Usage="EQUALITY">
        <Column Name="Column1" ColumnId="14" />
      </ColumnGroup>
      <ColumnGroup Usage="INEQUALITY">
        <Column Name="Column2" ColumnId="17" />
        <Column Name="Column3" ColumnId="21" />
      </ColumnGroup>
      <ColumnGroup Usage="INCLUDE">
        <Column Name="Column4" ColumnId="11" />
      </ColumnGroup>
    </MissingIndex>
  </MissingIndexGroup>
</MissingIndexes>
```

You need to create the missing index. Which code segment should you execute?

- A. CREATE NONCLUSTERED INDEX IX1 on Schema1.Table1 (Column1) INCLUDE (Column4) WHERE Column2 <> Column3
- B. CREATE NONCLUSTERED INDEX IX1 on Schema1.Table1 (Column1)
- C. CREATE NONCLUSTERED INDEX IX1 on Schema1.Table1 (Column1, Column2, Column3) INCLUDE (Column4)
- D. CREATE NONCLUSTERED INDEX IX1 on schema1.Table1 (Column1) INCLUDE (Column4)

Answer: C

NEW QUESTION 87

- (Exam Topic 8)

You create a disk-based table that contains the following script:

```
CREATE TABLE dbo.Products
(
  ProductId bigint IDENTITY(1,1),
  Name nvarchar(50) NULL,
  Description nvarchar(max) NULL,
  SKU char(10) NULL,
  CONSTRAINT PK_Products PRIMARY KEY CLUSTERED (ProductId)
) ON [PRIMARY]
GO
```

You need to prevent duplicate values in the SKU field. Which five code segments should you use?

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 |
|--------------------------|-------------|
| (SKU) | |
| ALTER TABLE dbo.Products | |
| CHECK | |
| FOREIGN KEY | |
| UNIQUE | |
| ADD CONSTRAINT | |
| CK_SKU | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: ALTER TABLE dbo.Products Box 2: ADD CONSTRAINT

Box 3: CK_SKU

Box 4: UNIQUE

Box 5: (SKU)

Note: The SQL command is:

ALTER TABLE <tablename> ADD CONSTRAINT

<constraintname> UNIQUE (

<columnname>

)

NEW QUESTION 88

- (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 91

- (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.

```
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:

NEW QUESTION 93

- (Exam Topic 8)

You plan to deploy two stored procedures name USP_1 and USP_2 that read data from a database. Your company identifies the following requirements for each stored procedure:

You need to identify which isolation level you must set for each stored procedure. The solution must minimize the number of locks.

Which isolation level should you identify?

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

| Isolation Levels | Answer Area |
|------------------|-----------------------|
| read committed | USP_1 Isolation level |
| read uncommitted | USP_2 Isolation level |
| repeatable read | |
| serializable | |
| snapshot | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: read uncommitted

READ UNCOMMITTED is the least restrictive isolation level because it ignores locks placed by other transactions. Transactions executing under READ UNCOMMITTED can read modified data values that have not yet been committed by other transactions; these are called "dirty" reads.

Box 2: SERIALIZABLE

Places a range lock on the data set, preventing other users from updating or inserting rows into the data set until the transaction is complete. This is the most

restrictive of the four isolation levels. Because concurrency is lower, use this option only when necessary. This option has the same effect as setting HOLDLOCK on all tables in all SELECT statements in a transaction.

References: [https://msdn.microsoft.com/en-us/library/tcbchxcb\(v=vs.110\).aspx](https://msdn.microsoft.com/en-us/library/tcbchxcb(v=vs.110).aspx)

NEW QUESTION 95

- (Exam Topic 8)

You have a database for a mission-critical web application. The database is stored on a SQL Server 2012 instance and is the only database on the instance. The application generates all T-SQL statements dynamically and does not use stored procedures. You need to maximize the amount of memory available for data caching.

Which advanced server option should you modify?

- A. Optimize for Ad hoc Workloads
- B. Enable Contained Databases
- C. Allow Triggers to Fire Others
- D. Scan for Startup Procs

Answer: A

NEW QUESTION 98

- (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 103

- (Exam Topic 8)

You have the following stored procedure.

```
CREATE PROCEDURE GetFile
    @FileName nvarchar(512)
AS
SELECT *
FROM Files
WHERE FileName = @FileName
```

The stored procedure takes much longer to execute than expected.

While reviewing the execution plan of the stored procedure, you discover the following predicate for a Clustered Index Scan operator.

```
CONVERT_IMPLICIT(nvarchar(512), [SampleDatabase].[dbo].[Files].[FileName], 0)=[@1]
```

You need to resolve the performance issue. What should you do?

- A. Change the @FileName parameter from nvarchar(512) to varchar(512).
- B. Change the FileName column from varchar(512) to nvarchar(512).
- C. Add a NOLOCK query hint to the SELECT statement.
- D. Convert the table to a memory-optimized table.
- E. Add a FORCESEEK query hint to the SELECT statement.

Answer: A

Explanation:

When using a variable, make sure that the datatype matches the column's datatype. We suspect that the issue is that the variable is NVARCHAR (512) whilst the table column is VARCHAR (512). This is indicated by the CONVERT_IMPLICIT operator in the execution plan.

References:

https://sqlserverperformance.wordpress.com/2009/02/02/beware-of-convert_implicit-in-a-sql-execution-plan/

NEW QUESTION 104

- (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.)

| Functions | Answer Area |
|-----------------------------------|---|
| sys.dm_db_index_usage_stats | Identify unused indexes. |
| sys.dm_db_index_operational_stats | Identify which indexes should be created. |
| sys.dm_db_index_physical_stats | |
| sys.dm_db_missing_index_columns | |
| sys.dm_db_missing_index_details | |
| sys.dm_db_missing_index_groups | |

- 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 108

- (Exam Topic 8)

You plan to create a new table that has the following requirements:

- Uses a GUID data type as the primary key.
- Uses a clustered index as the primary key.
- Minimizes fragmentation.

You need to recommend which option to include in the CREATE statement. Which option should you include?

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

- A. NEWID
- B. @@IDENTITY
- C. NEWSEQUENTIALID
- D. IDENTITY

Answer: C

NEW QUESTION 110

- (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
SERIALIZABLE
WITH(UPDLOCK)
WITH(XLOCK)

CREATE PROCEDURE Proc1 (@ParamID int)
AS
SET TRANSACTION ISOLATION LEVEL
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 115

- (Exam Topic 8)

You have a database named Database1. Database1 has two stored procedures named Proc1 and Proc2 and a table named Table1. Table1 has millions of rows. Proc1 updates data in Table1. Proc2 reads data from Table1.

You discover that when Proc1 is executed to update more than 4,000 rows, Proc2 is blocked. The block affects all rows, including those that are not being updated by Proc1.

You need to ensure that when Proc1 is executing, Proc2 can access the data in Table1 that Proc1 is not updating.

What should you change Proc1 to do?

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

- A. Update less than 4,000 rows simultaneously.
- B. Use the PAGLOCK table hint.
- C. Wait for Proc2 to complete.
- D. Use the ROWLOCK table hint.

Answer: A

NEW QUESTION 118

- (Exam Topic 8)

You execute the following code:

```
CREATE TABLE UserInfo
(
  ID int NOT NULL IDENTITY (1, 1)
  CONSTRAINT PK_UserInfo PRIMARY KEY CLUSTERED,
  UserName varchar(100) NOT NULL,
  Manager varchar(100) NULL,
  HireDate date NOT NULL,
  PerformanceReviewScore int NULL
);
```

You have a stored procedure that includes the following SELECT statement:

```
SELECT UserName, PerformanceReviewScore
FROM UserInfo
WHERE Manager = 'Ben Smith';
```

You need to create a covering index on UserInfo. Which code segment should you execute?

- A. CREATE NONCLUSTERED INDEX [IX_Covering_Index] ON UserInfo
 (
 [Manager] ASC
);
- B. CREATE NONCLUSTERED INDEX [IX_Covering_Index] ON UserInfo
 (
 [UserName] ASC,
 [PerformanceReviewScore] ASC,
);
- C. CREATE NONCLUSTERED INDEX [IX_Covering_Index] ON UserInfo
 (
 [Manager] ASC,
 [PerformanceReviewScore] ASC,
 [UserName] ASC
);
- D. CREATE NONCLUSTERED INDEX [IX_Covering_Index] ON UserInfo
 (
 [UserName] ASC,
 [Manager] ASC
);

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

Answer: C

NEW QUESTION 120

- (Exam Topic 8)

You create a table named Customers by using the following code segment:

```
CREATE TABLE dbo.Customers
(
  id int primary key,
  name char(10)
)
```

You create a non-clustered index named IX_Name on the name column.

You write the following query to retrieve all of the customers that have a name that starts with the letters SMI:

```
SELECT * FROM dbo.Customers
WHERE 'smi' = LEFT(name, 3)
```

You discover that the query performs a table scan. You need to ensure that the query uses the index. What should you do?

- A. Replace LEFT(name,3) = 'smi' by using name like 'smi%'
- B. Replace LEFT(name,3) = 'smi' by using substring(name,1,3) = 'smi'
- C. Recreate IX_Name as a unique index
- D. Recreate IX Name as a clustered index

Answer: A

NEW QUESTION 123

- (Exam Topic 8)

You execute the following code.

```
CREATE TABLE HumanResources.Employees
(
    EmployeeID int IDENTITY(1,1) PRIMARY KEY,
    ContactID int NOT NULL
        FOREIGN KEY REFERENCES Person.Contact(ContactID),
    JobTitle varchar(100)
);
GO

CREATE INDEX IX_Employees
ON HumanResources.Employee(JobTitle);
GO
```

After populating the Employees table with 10,000 rows, you execute the following query:

```
SELECT EmployeeID, JobTitle
FROM HumanResources.Employee
WHERE SUBSTRING(JobTitle,1,1) = 'C'
```

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 JobTitle column for the partition scheme.
- B. Change SUBSTRING(JobTitle,1,1) = 'C' to JobTitle LIKE 'c%'
- C. Change SUBSTRING(JobTitle, 1,1) = 'c' to LEFT(JobTitle,1) = 'c'.
- D. Replace IX_Employees with a clustered index.

Answer: B

NEW QUESTION 125

- (Exam Topic 8)

You have a SQL Server 2012 database named DB1 that is accessed by 650 concurrent users.

You need to log all of the queries to DB1 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 SQL Server Agent job that retrieves information from the sys.dm_tran_session_transactions dynamic management views
- C. A SQL Server Agent job that retrieves information from the sys.dm_tran_active_transactions dynamic management views
- D. A script that enables trace flags

Answer: A

NEW QUESTION 128

- (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 129

- (Exam Topic 8)

You have a Microsoft SQL Azure database. You have the following stored procedure:

```
01 CREATE PROC up_employees
02     @ID int,
03     @Name nvarchar(50)
04 AS
05
06 SELECT Name AS OriginalName
07 FROM HR.Employees
08
09 WHERE ID = @ID;
10
11 UPDATE HR.Employees
12 SET Name = @Name
13
14 WHERE ID = @ID;
```

You discover that the stored procedure periodically fails to update HR.Employees. You need to ensure that HR.Employees is always updated when up_employees executes. The solution must minimize the amount of time required for the stored procedure to execute and the number of locks held. What should you do?

- A. Add the following line of code to line 05: SET TRANSACTION ISOLATION LEVEL SNAPSHOT
- B. Add the following line of code to line 13: WITH (UPDLOCK)
- C. Add the following line of code to line 05: SET TRANSACTION ISOLATION LEVEL SERIALIZABLE
- D. Add the following line of code to line 08: WITH (UPDLOCK)

Answer: D

NEW QUESTION 134

- (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 135

- (Exam Topic 8)

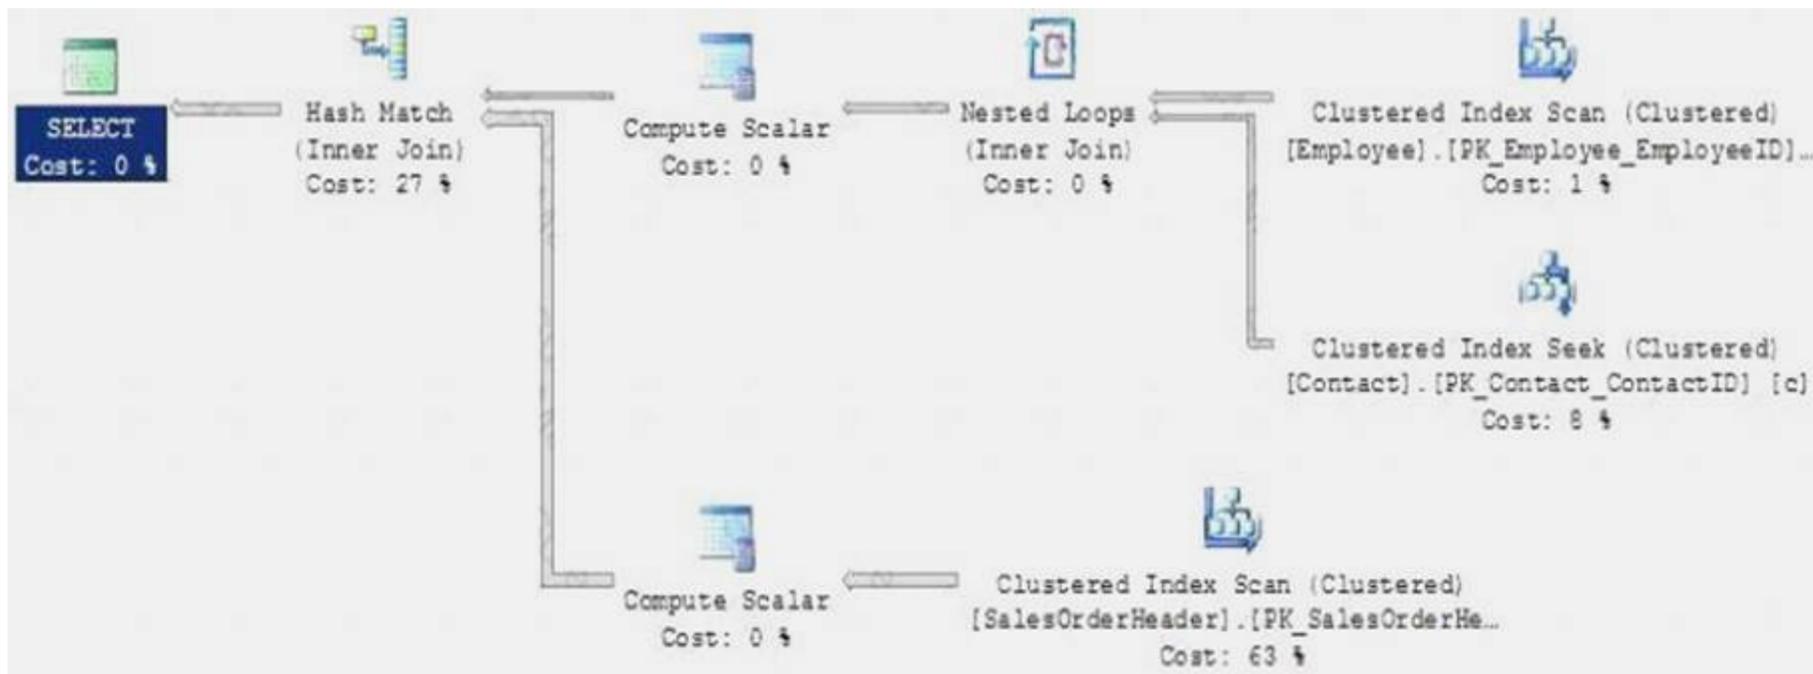
You have a database that contains three tables. The tables are configured as shown in the following table.

| Table | Primary key index |
|------------------|----------------------------------|
| SalesOrderHeader | PK_SalesOrderHeader_SalesOrderID |
| Employee | PK_Employee_EmployeeID |
| Contact | PK_Contact_ContactID |

You have the following query:

```
SELECT soh.SalesPersonID,
       c.FirstName + ' ' + c.LastName AS FullName,
       c.EmailAddress,
       e.Title,
       soh.SubTotal,
       YEAR(soh.OrderDate) AS Year
FROM SalesOrderHeader soh
INNER JOIN Employee e
    ON soh.SalesPersonID = e.EmployeeID
INNER JOIN Contact c
    ON e.ContactID = c.ContactID
WHERE soh.OrderDate >= '1/1/2012'
```

The execution plan for the query is shown in the exhibit. (Click the Exhibit button.)



You need to create one index to minimize the amount of time it takes to execute the query.

What should you do?

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

| Columns | Answer Area |
|--------------------------------|-----------------------------------|
| Contact.EmailAddress | Indexed Columns Column |
| Contact.FirstName | Included Columns Column |
| Contact.LastName | Column |
| Employee.Title | Column |
| SalesOrderHeader.OrderDate | |
| SalesOrderHeader.SalesPersonID | |
| SalesOrderHeader.SubTotal | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

Covering index: A type of index that includes all the columns that are needed to process a particular query. For example, your query might retrieve the FirstName and LastName columns from a table, based on a value in the ContactID column. You can create a covering index that includes all three columns.

NEW QUESTION 137

- (Exam Topic 8)

You use SQL Server 2012 to store data used by an e-commerce application.

You develop a stored procedure named sp1. Sp1 is used to read the price of all the products sold on the e-commerce site.

You need to ensure that sp1 can read data even while another transaction is modifying the price of a product. Sp1 must only read committed data.

Which transaction isolation level should you use in sp1?

- A. Serializable
- B. Snapshot
- C. Repeatable read
- D. Read committed

Answer: B

NEW QUESTION 139

- (Exam Topic 8)

You are designing a new database table that will be used for reporting. You define the table by using the following statement.

```
CREATE TABLE Customers
(CustomerId BIGINT,
LastName VARCHAR(255),
FirstName VARCHAR(255),
SalesPersonId INT,
Address1 VARCHAR(255),
Address2 VARCHAR(255),
City VARCHAR(255),
State CHAR(2),
PostalCode VARCHAR(10),
Country CHAR(2))
```

You need to store the data in the table by using the least amount of storage space possible. Which storage option should you use?

- A. a clustered index
- B. a clustered columnstore index
- C. a nonclustered index
- D. In-Memory OLTP

Answer: B

Explanation:

Columnstore indexes work well for mostly read-only queries that perform analysis on large data sets. This would fit this scenario as the table will be used for reporting.

Columnstore Index benefits include Columnstore Index benefits high compression rates, which improve query performance by using a smaller in-memory footprint. In turn, query performance can improve because SQL Server can perform more query and data operations in-memory.

Use the columnstore index to achieve up to 10x query performance gains over traditional row-oriented storage, and up to 7x data compression over the uncompressed data size.

References: [https://msdn.microsoft.com/en-us/library/gg492088\(v=sql.120\).aspx](https://msdn.microsoft.com/en-us/library/gg492088(v=sql.120).aspx)

NEW QUESTION 143

- (Exam Topic 8)

You run the following code:

```
CREATE TABLE dbo.Orders
(
    Id int CONSTRAINT PK_Order_Id PRIMARY KEY,
    Amount decimal,
    Details xml
);
```

You need to ensure that the root node of the XML data stored in the Details column is <Order_Details>. What should you implement?

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

- A. A user-defined data type
- B. An XML index
- C. A Data Definition Language (DDL) trigger
- D. A data manipulation language (DML) trigger
- E. An XML schema collection

Answer: E

Explanation:

<http://msdn.microsoft.com/en-us/library/ms187856.aspx>

NEW QUESTION 146

- (Exam Topic 8)

You have an index for a table in a SQL Azure database. The database is used for Online Transaction Processing (OLTP).

You discover that many page splits occur when records are inserted or updated in the table. You need to minimize the number of page splits.

What should you set from the index options?

- A. FILLFACTOR = 0
- B. STATISTICS_NORECOMPUTE = OFF
- C. STATISTICS_NORECOMPUTE = ON
- D. FILLFACTOR = 80

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms188783.aspx> <http://msdn.microsoft.com/en-us/library/ms177459.aspx>

NEW QUESTION 148

- (Exam Topic 8)

You need to identify which long running transactions use an index. Which dynamic management view should you use?

- A. sys.dm_exec_query_optimizer_info
- B. sys.dm_exec_connections
- C. sys.dm_exec_query_stats
- D. sys.dm_exec_sessions

Answer: A

NEW QUESTION 151

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

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 First Try
- * 70-464 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 70-464 Practice Test Here](#)