



Vendor: Microsoft

Exam Code: 70-464

Exam Name: Developing Microsoft SQL Server 2014
Databases Exam

Version: DEMO

Case Study 1: Scenario 1 (Question 47 ~ Question 57)

Case Study 2: Scenario 2 (Question 58 ~ Question 68)

Case Study 3: Scenario 3 (Question 69 ~ Question 79)

Case Study 4: Scenario 4 (Question 80 ~ Question 85)

Case Study 5: Litware Inc (Question 86 ~ Question 96)

Case Study 6: Coho Winery (Question 97 ~ Question 108)

Case Study 7: Fourth Coffee (Question 156 ~ Question 165)

Background

Corporate Information

Fourth Coffee is global restaurant chain. There are more than 5,000 locations worldwide.

Physical Locations

Currently a server at each location hosts a SQL Server 2012 instance. Each instance contains a database called StoreTransactions that stores all transactions from point of sale and uploads summary batches nightly.

Each server belongs to the COFFECORP domain. Local computer accounts access the StoreTransactions database at each store using sysadmin and datareaderwriter roles.

QUESTION 1

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?

- ☐ A. SET XACT_ABORT ON
BEGIN TRY
BEGIN TRANSACTION
...Statement 1
...Statement 2
COMMIT TRANSACTION
END TRY
BEGIN CATCH
ROLLBACK TRANSACTION
END CATCH
- ☐ B. SET XACT_ABORT OFF
BEGIN TRY
...Statement 1
END TRY
BEGIN TRY
...Statement 2
END TRY
BEGIN CATCH
THROW
END CATCH
- ☐ C. SET XACT_ABORT ON
BEGIN TRANSACTION
...Statement 1
...Statement 2
If @@ERROR <> 0
ROLLBACK
ELSE
COMMIT TRANSACTION
- ☐ 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

- 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.

QUESTION 2

Drag and Drop Question

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

```
CREATE TABLE Product (  
ProductID int IDENTITY(1,1) PRIMARY KEY,  
ProductName varchar(30) NOT NULL,  
LastUpdatedDate smalldatetime,  
LastUpdatedBy varchar(128))
```

```
CREATE TABLE ProductAudit (  
ProductAuditID int IDENTITY(1,1) PRIMARY KEY,  
OldProductID int NOT NULL,  
OldProductName varchar(30) NOT NULL,  
UpdatedDate smalldatetime,  
UpdatedBy varchar(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</pre>	
<pre>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</pre>	
<pre>END</pre>	
<pre>COMMIT TRANSACTION</pre>	
<pre>IF @@ERROR <> 0 ROLLBACK</pre>	
<pre>CREATE TRIGGER ProductAuditTrigger ON Product FOR INSERT, UPDATE, DELETE AS BEGIN</pre>	

Answer:

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 </pre>	<pre> CREATE TRIGGER ProductAuditTrigger ON Product FOR INSERT, UPDATE, DELETE AS BEGIN </pre>
<pre> 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 </pre>	<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 </pre>
<pre> END </pre>	<pre> COMMIT TRANSACTION </pre>
<pre> COMMIT TRANSACTION </pre>	<pre> IF @@ERROR <> 0 ROLLBACK </pre>
<pre> IF @@ERROR <> 0 ROLLBACK </pre>	<pre> END </pre>
<pre> CREATE TRIGGER ProductAuditTrigger ON Product FOR INSERT, UPDATE, DELETE AS BEGIN </pre>	

QUESTION 3

You administer an instance of SQL Server 2014.

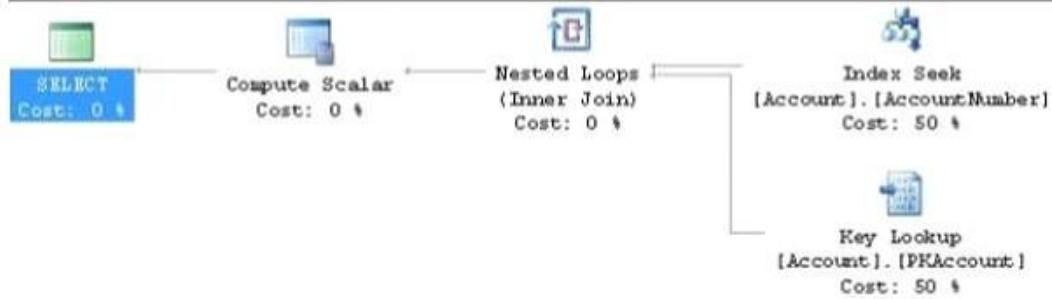
You are tasked with tuning a common set of queries.

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'
```

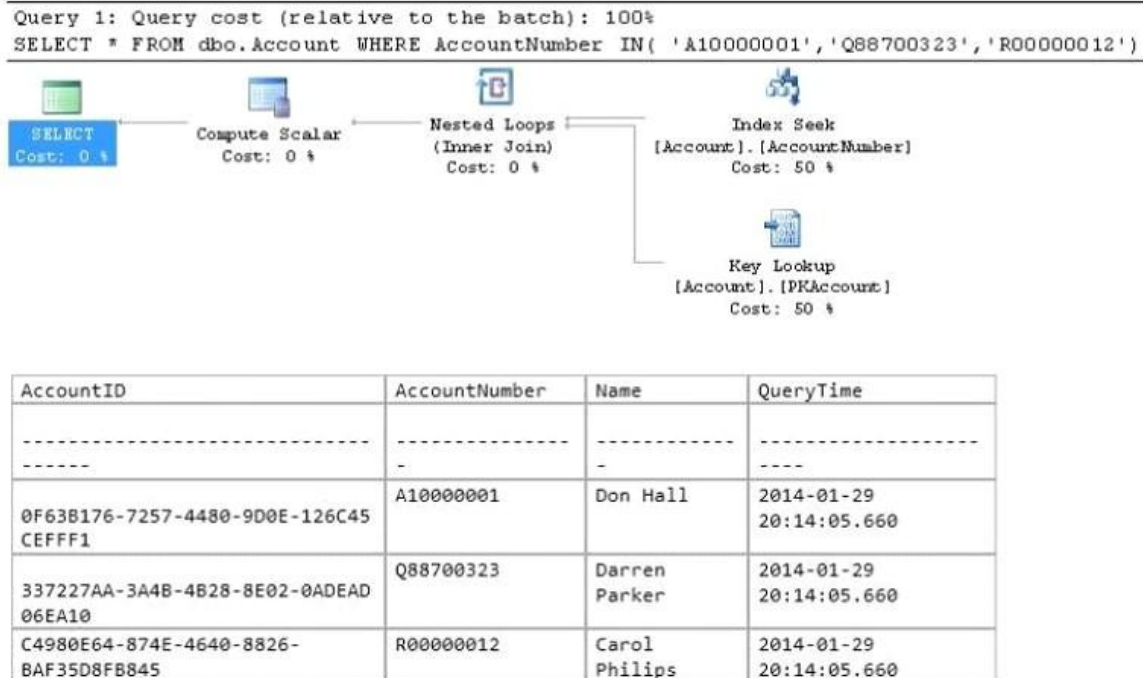
Query 1: Query cost (relative to the batch): 100%

SELECT * FROM [dbo].[Account] WHERE [AccountNumber]=01



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
```



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 Query Time 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

QUESTION 4

You use SQL Server 2014. The physical server is a dedicated database server that has 120GB of RAM available. There is approximately 50GB of storage space available on a slow local disk.
You create a new stored procedure.

You decide you need to temporarily hold approximately 300,000 rows from two tables, from which you will compute two complex business scores.

The stored procedure will use temporary storage defined as follows:

```
AccountNumber char(10) NOT NULL
YearToDateSalesTotal decimal(15,2) NULL
SalesScore int NULL
FutureSalesExpectationScore int NULL
```


The code will make several passes through the data, applying complex calculations before writing the data to a permanent disk-based table in the same database from which it reads the data. For this stored procedure, you need to deal with temporary data in the most efficient way to minimize physical disk pressure.

What should you do?

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

- ☐ A. `CREATE TYPE dbo.AccountScoringModel as TABLE`
(
 AccountNumber char(10) COLLATE Latin1_General_100_BIN2 NOT NULL ,
 YearToDateSalesTotal decimal(15,2) NULL,
 SalesScore int NULL,
 FutureSalesExpectationScore int NULL,
 INDEX AccountNumber HASH (AccountNumber) WITH (BUCKET_COUNT = 25000)
) WITH (MEMORY_OPTIMIZED = ON)
GO
DECLARE @AccountScoring as dbo.AccountScoringModel
- ☐ B. `DECLARE @AccountScoring as TABLE`
(
 AccountNumber char(10) NOT NULL,
 YearToDateSalesTotal decimal(15,2) NULL,
 SalesScore int NULL,
 FutureSalesExpectationScore int NULL
)
- ☐ C. `CREATE TABLE #AccountScoring`
(
 AccountNumber char(10) NOT NULL,
 YearToDateSalesTotal decimal(15,2) NULL,
 SalesScore int NULL,
 FutureSalesExpectationScore int NULL
)
- ☐ D. `CREATE TYPE dbo.AccountScoringModel as TABLE`
(
 AccountNumber char(10) COLLATE Latin1_General_100_BIN2 NOT NULL ,
 YearToDateSalesTotal decimal(15,2) NULL,
 SalesScore int NULL,
 FutureSalesExpectationScore int NULL,
 INDEX AccountNumber HASH (AccountNumber) WITH (BUCKET_COUNT = 120)
) WITH (MEMORY_OPTIMIZED = ON)
GO
DECLARE @AccountScoring as dbo.AccountScoringModel

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

Answer: A

Explanation:

* 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. If the index key contains a lot of duplicate values, on average there are more than 10 rows for each index key value, use a nonclustered index instead

You may not always be able to predict how many values a particular index key may have or will have. Performance should be acceptable if the BUCKET_COUNT value is within 5 times of the actual number of key values.

QUESTION 5

You are creating a table to support an application that will cache data outside of SQL Server. The application will detect whether cached values were changed before it updates the values. You need to create the table, and then verify that you can insert a row into the table. Which code segment should you use?

- ☐ A.

```
CREATE TABLE Table1
(
    ID int IDENTITY(1,1),
    Name varchar(100),
    Version uniqueidentifier DEFAULT (NEWID())
)
INSERT INTO Table1 (Name, Version)
VALUES ('Smith, Ben', NEWID())
```
- ☐ B.

```
CREATE TABLE Table1
(
    ID int IDENTITY(1,1),
    Name varchar(100),
    Version uniqueidentifier DEFAULT (NEWID())
)
INSERT INTO Table1 (Name)
VALUES ('Smith, Ben')
```
- ☐ C.

```
CREATE TABLE Table1
(
    ID int IDENTITY(1,1),
    Name varchar(100),
    Version rowversion
)
INSERT INTO Table1 (Name)
VALUES ('Smith, Ben')
```
- ☐ D.

```
CREATE TABLE Table1
(
    ID int IDENTITY(1,1),
    Name varchar(100),
    Version rowversion
)
INSERT INTO Table1 (Name, Version)
VALUES ('Smith, Ben', NEWID())
```

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

Answer: C

Explanation:

<http://msdn.microsoft.com/en-us/library/ms182776.aspx>

<http://msdn.microsoft.com/en-us/library/ms187942.aspx>

<http://msdn.microsoft.com/en-us/library/ms190348.aspx>

QUESTION 6

Your company has a SQL Azure subscription.

You implement a database named Database1. Database1 has two tables named Table1 and Table2.

You create a stored procedure named sp1. Sp1 reads data from Table1 and inserts data into Table2. A user named User1 informs you that he is unable to run sp1.

You verify that User1 has the SELECT permission on Table1 and Table2.

You need to ensure that User1 can run sp1. The solution must minimize the number of permissions assigned to User1.

What should you do?

- A. Grant User1 the INSERT permission on Table2.
- B. Add User1 to the db_datawriter role.
- C. Change sp1 to run as the sa user.
- D. Grant User1 the EXECUTE permission on sp1.

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms191291.aspx>

QUESTION 7

You have a SQL Server 2012 database named DB1.

You have a backup device named Device1.

You discover that the log file for the database is full.

You need to ensure that DB1 can complete transactions. The solution must not affect the chain of log sequence numbers (LSNs).

Which code segment should you execute?

- A. BACKUP LOG DB1 TO Device1 WITH COPY_ONLY
- B. BACKUP LOG DB1 TO Device1
- C. BACKUP LOG DB1 TO Device1 WITH NCRECCVERY
- D. BACKUP LOG D31 TO Device1 WITH TRUNCATE ONLY

Answer: B

Explanation:

<http://msdn.microsoft.com/en-us/library/ms186865.aspx>

<http://msdn.microsoft.com/en-us/library/ms179478.aspx>

<http://msdn.microsoft.com/en-us/library/ms190925.aspx>

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives®

10% Discount Coupon Code: ASTR14