

Exam Questions 70-762

Developing SQL Databases (beta)

<https://www.2passeasy.com/dumps/70-762/>



NEW QUESTION 1

You use Microsoft SQL Server Profiler to evaluate a query named Query1. The Profiler report indicates the following issues:

- At each level of the query plan, a low total number of rows are processed.
- The query uses many operations. This results in a high overall cost for the query. You need to identify the information that will be useful for the optimizer. What should you do?

- A. Start a SQL Server Profiler trace for the event class Auto Stats in the Performance event category.
- B. Create one Extended Events session with the sqlserver.missing_column_statistics event added.
- C. Start a SQL Server Profiler trace for the event class Soft Warnings in the Errors and Warnings event category.
- D. Create one Extended Events session with the sqlserver.missing_join_predicate event added.

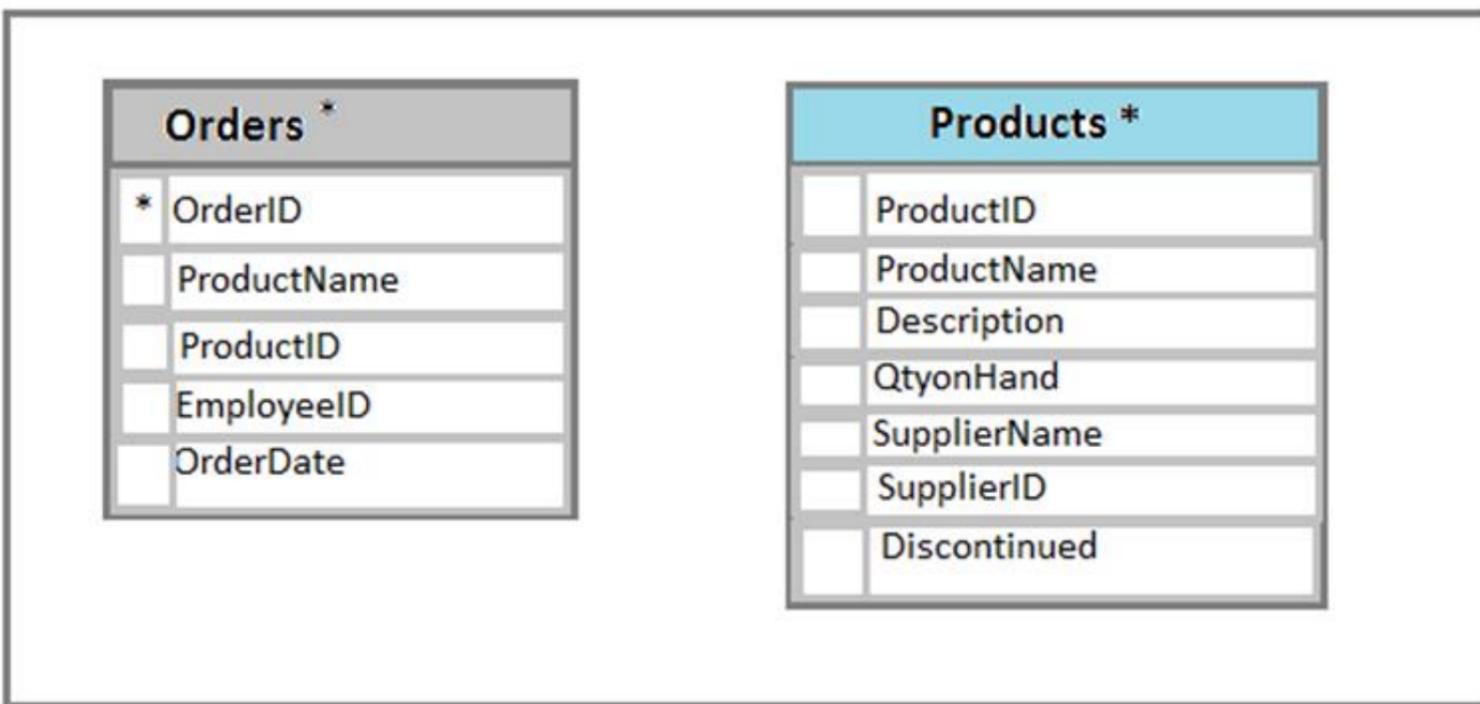
Answer: D

Explanation: The Missing JoinPredicate event class indicates that a query is being executed that has no join predicate. This could result in a long-running query.

NEW QUESTION 2

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

You have a database named Sales that contains the following database tables: Customer, Order, and Products. The Products table and the Order table are shown in the following diagram.



The customer table includes a column that stores the data for the last order that the customer placed.

You plan to create a table named Leads. The Leads table is expected to contain approximately 20,000 records. Storage requirements for the Leads table must be minimized.

The Leads table must include the columns described in the following table.

Column name	Description
LeadID	This column stores a unique value that is automatically assigned for each lead.
IsCustomer	This column indicates whether the lead is for a current customer.

The data types chosen must consume the least amount of storage possible. You need to select the appropriate data types for the Leads table. In the table below, identify the data type that must be used for each table column. NOTE: Make only one selection in each column.

Answer Area

Data type	LeadID	IsCustomer
smallint	<input type="radio"/>	<input type="radio"/>
int	<input type="radio"/>	<input type="radio"/>
binary	<input type="radio"/>	<input type="radio"/>
numeric	<input type="radio"/>	<input type="radio"/>
bit	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation: Bit is a Transact-SQL integer data type that can take a value of 1, 0, or NULL.

Smallint is a Transact-SQL integer data type that can take a value in the range from -32,768 to 32,767. int, bigint, smallint, and tinyint (Transact-SQL) Exact-number data types that use integer data.

Data type	Range	Storage
bigint	-2 ⁶³ (-9,223,372,036,854,775,808) to 2 ⁶³ -1 (9,223,372,036,854,775,807)	8 Bytes
int	-2 ³¹ (-2,147,483,648) to 2 ³¹ -1 (2,147,483,647)	4 Bytes
smallint	-2 ¹⁵ (-32,768) to 2 ¹⁵ -1 (32,767)	2 Bytes
tinyint	0 to 255	1 Byte

References: <https://msdn.microsoft.com/en-us/library/ms187745.aspx> <https://msdn.microsoft.com/en-us/library/ms177603.aspx>

NEW QUESTION 3

You are analyzing the performance of a database environment.

Applications that access the database are experiencing locks that are held for a large amount of time. You are experiencing isolation phenomena such as dirty, nonrepeatable and phantom reads.

You need to identify the impact of specific transaction isolation levels on the concurrency and consistency of data.

What are the consistency and concurrency implications of each transaction isolation level? To answer, drag the appropriate isolation levels to the correct locations. Each isolation level 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.

Isolation levels

read committed

serializable

read uncommitted

repeatable read

Concurrency

Isolation Level

Isolation level

Isolation level

Isolation level

Isolation level

Consistency

Answer:

Explanation: Read Uncommitted (aka dirty read): A transaction T1 executing under this isolation level can access data changed by concurrent transaction(s).

Pros: No read locks needed to read data (i.e. no reader/writer blocking). Note, T1 still takes transaction duration locks for any data modified.

Cons: Data is not guaranteed to be transactionally consistent.

Read Committed: A transaction T1 executing under this isolation level can only access committed data. Pros: Good compromise between concurrency and consistency.

Cons: Locking and blocking. The data can change when accessed multiple times within the same transaction.

Repeatable Read: A transaction T1 executing under this isolation level can only access committed data with an additional guarantee that any data read cannot change (i.e. it is repeatable) for the duration of the transaction.

Pros: Higher data consistency.

Cons: Locking and blocking. The S locks are held for the duration of the transaction that can lower the concurrency. It does not protect against phantom rows.

Serializable: A transaction T1 executing under this isolation level provides the highest data consistency including elimination of phantoms but at the cost of reduced concurrency. It prevents phantoms by taking a range lock or table level lock if range lock can't be acquired (i.e. no index on the predicate column) for the duration of the transaction.

Pros: Full data consistency including phantom protection.

Cons: Locking and blocking. The S locks are held for the duration of the transaction that can lower the concurrency.

References:

<https://blogs.msdn.microsoft.com/sqlcat/2011/02/20/concurrency-series-basics-of-transaction-isolation-levels/>

NEW QUESTION 4

Note: This question is part of a series of questions that present the same scenario. Each question in this series contains a unique solution. Determine whether the solution meets the stated goals.

The Account table was created by using the following Transact-SQL statement:

```
CREATE TABLE Account
(
    AccountNumber int NOT NULL,
    ProductCode char(2) NOT NULL,
    Status tinyint NOT NULL,
    OpenDate date NOT NULL,
    CloseDate date,
    Balance decimal(15,2),
    AvailableBalance decimal(15,2)
);
```

There are more than 1 billion records in the Account table. The Account Number column uniquely identifies each account. The ProductCode column has 100 different values. The values are evenly distributed in the table. Table statistics are refreshed and up to date. You frequently run the following Transact-SQL SELECT statements:

```
SELECT ProductCode, SUM(Balance) AS TotalSUM FROM Account WHERE ProductCode
<> 'CD' GROUP BY ProductCode;
SELECT AccountNumber, Balance FROM Account WHERE Production = 'CD'
```

You must avoid table scans when you run the queries. You need to create one or more indexes for the table. Solution: You run the following Transact-SQL statement:

```
CREATE CLUSTERED INDEX PK_Account On Account(AccountNumber);
CREATE NONCLUSTERED INDEX IX_Account_ProductCode On Account(ProductCode)
INCLUDE (Balance);
```

Does the solution meet the goal?

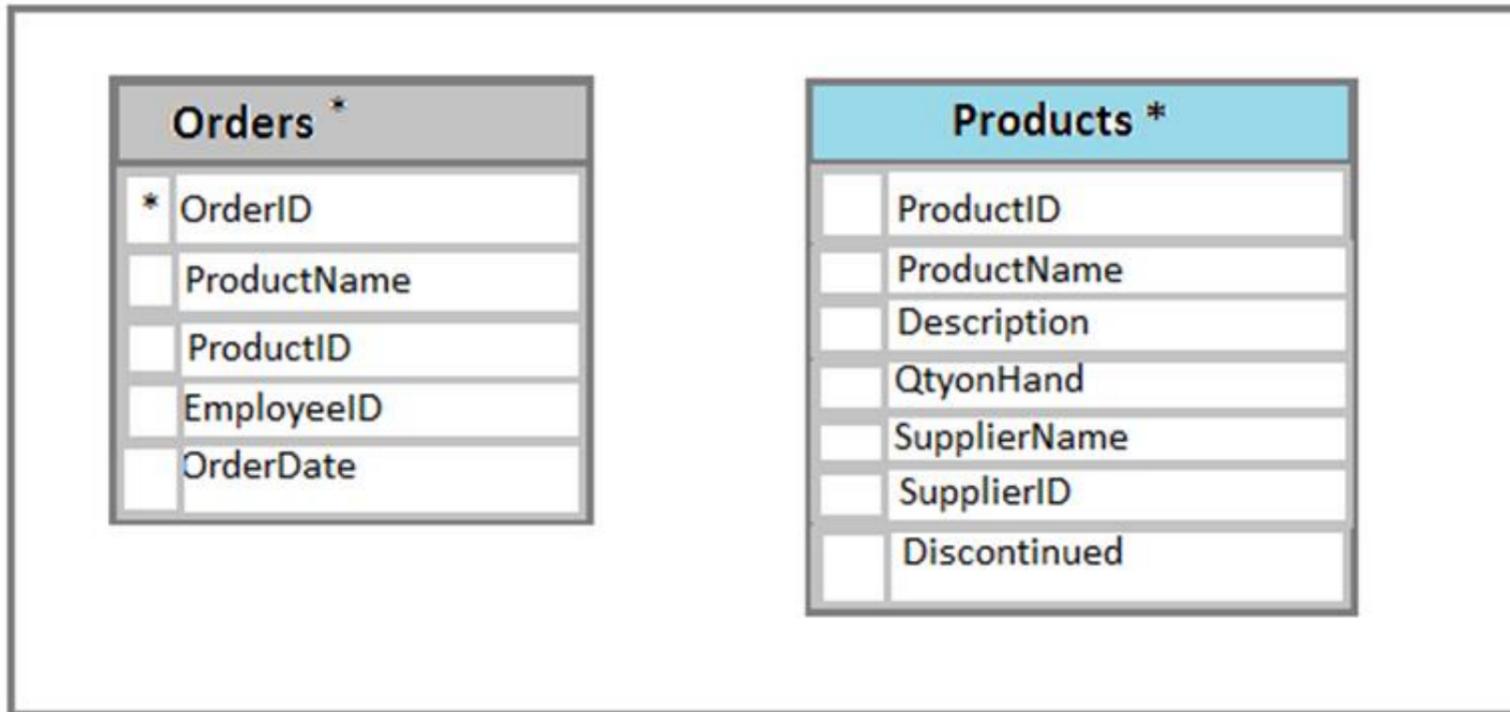
- A. Yes
- B. No

Answer: A

Explanation: Create a clustered index on the AccountNumber column as it is unique. Create a nonclustered index that includes the ProductCode column. References: <https://msdn.microsoft.com/en-us/library/ms190457.aspx>

NEW QUESTION 5

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in the series. You have a database named Sales that contains the following database tables. Customer, Order, and Products. The Products table and the order table shown in the following diagram.



The Customer table includes a column that stores the date for the last order that the customer placed. You plan to create a table named Leads. The Leads table is expected to contain approximately 20,000 records. Storage requirements for the Leads table must be minimized. You need to begin to modify the table design to adhere to third normal form. Which column should you remove for each table? To answer? drag the appropriate column names to the correct locations. Each column name 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.

Columns

- ProductID
- ProductName
- Description
- EmployeeID
- OrderDate
- SupplierName
- SupplierID
- Discontinued

•••••

Answer Area

	Table	Column to remove
Products		Column
Orders		Column

Answer:

Explanation: In the Products table the SupplierName is dependent on the SupplierID, not on the ProductID. In the Orders table the ProductName is dependent on the ProductID, not on the OrderID. Note:

A table is in third normal form when the following conditions are met:

- * It is in second normal form.
- * All non-primary fields are dependent on the primary key.

Second normal form states that it should meet all the rules for First 1Normal Form and there must be no partial dependencies of any of the columns on the primary key.

First normal form (1NF) sets the very basic rules for an organized database:

- * Define the data items required, because they become the columns in a table. Place related data items in a table.
- * Ensure that there are no repeating groups of data.
- * Ensure that there is a primary key.

References: <https://www.tutorialspoint.com/sql/third-normal-form.htm>

NEW QUESTION 6

Background

You have a database named HR1 that includes a table named Employee.

You have several read-only, historical reports that contain regularly changing totals. The reports use multiple queries to estimate payroll expenses. The queries run concurrently. Users report that the payroll estimate reports do not always run. You must monitor the database to identify issues that prevent the reports from running.

You plan to deploy the application to a database server that supports other applications. You must minimize the amount of storage that the database requires.

Employee Table

You use the following Transact-SQL statements to create, configure, and populate the Employee table:

```
CREATE TABLE dbo.Employee
    (EmployeeId INT PRIMARY KEY,
    LastName varchar(50),
    FirstName varchar(50),
    DepartmentId int,
    HireDate datetime,
    TerminationDate datetime,
    SupervisorId int,
    CostCenterNumber int,
    EmployeeStatus int,
    EmployeePayRate int)
) GO
CREATE INDEX IX_1 on dbo.Employee (LastName, FirstName, DepartmentId) INCLUDE (HireDate)
CREATE INDEX IX_2 on dbo.Employee (LastName) INCLUDE (EmployeeId, FirstName, DepartmentId)
CREATE INDEX IX_3 on dbo.Employee (LastName, FirstName) INCLUDE (DepartmentId)
CREATE INDEX IX_4 on dbo.Employee (LastName, FirstName) INCLUDE (HireDate, DepartmentId)
GO

INSERT INTO Employee (EmployeeID, LastName, CostCenterNumber) VALUES(1001, 'Employee A', 3001001)
INSERT INTO Employee (EmployeeID, LastName, CostCenterNumber) VALUES(1002, 'Employee B', 3001001)
GO
```

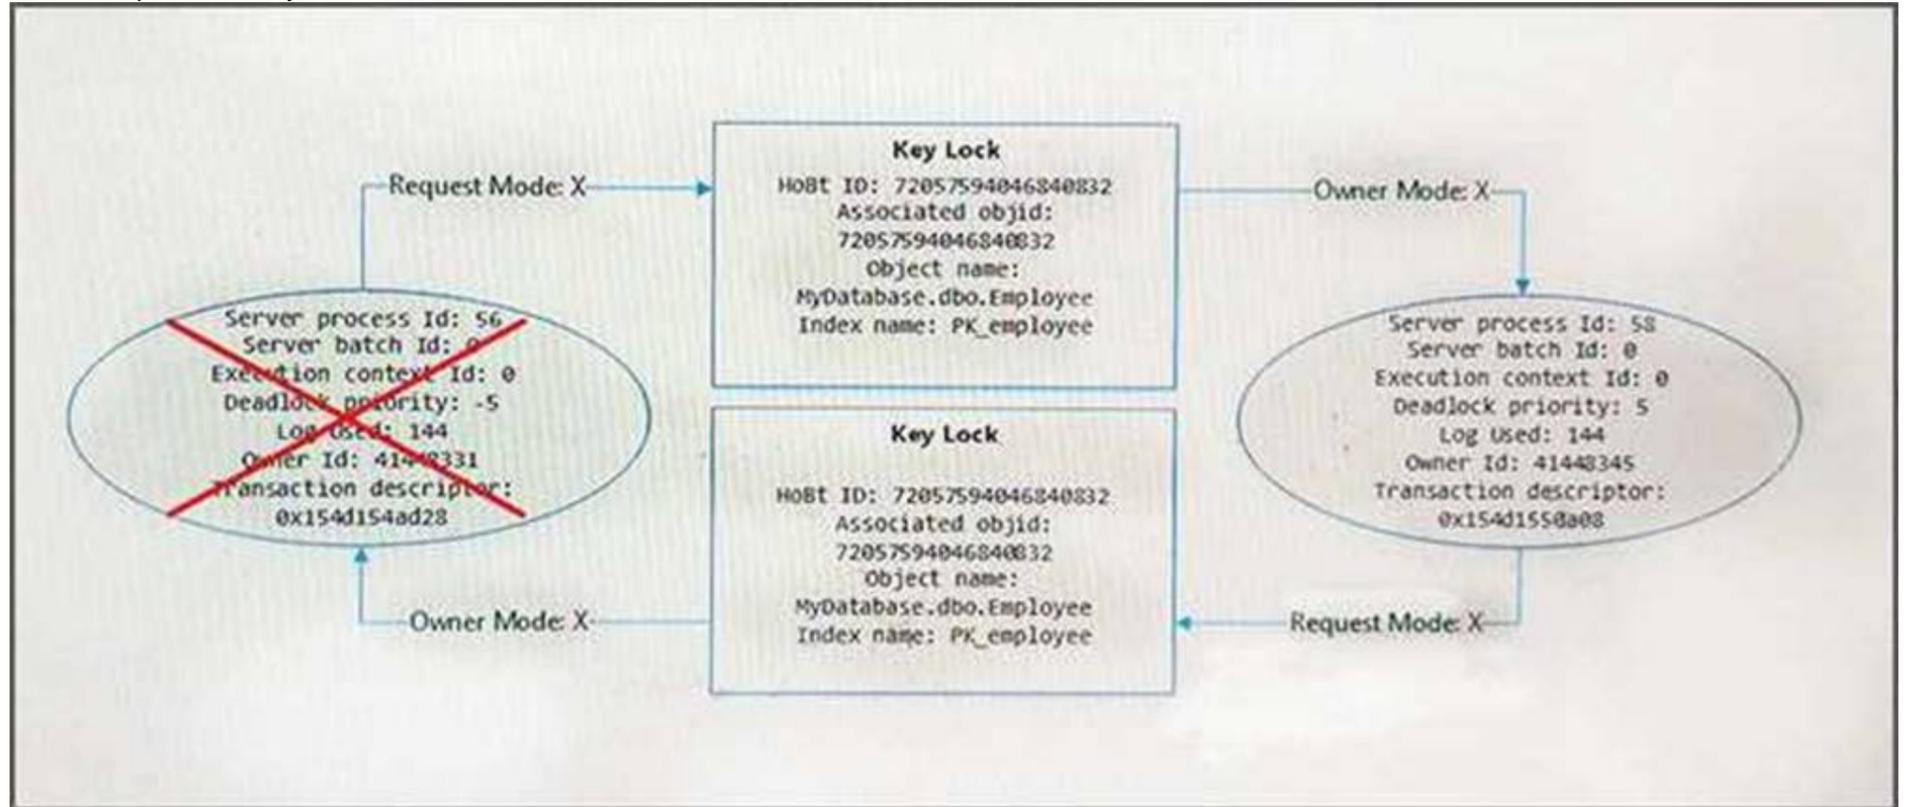
Application

You have an application that updates the Employees table. The application calls the following stored procedures simultaneously and asynchronously:

- UspA: This stored procedure updates only the EmployeeStatus column.
- UspB: This stored procedure updates only the EmployeePayRate column.

The application uses views to control access to data. Views must meet the following requirements:

- Allow user access to all columns in the tables that the view accesses.
- Restrict updates to only the rows that the view returns. Exhibit



You view the Deadlock Graph as shown in the exhibit. (Click the Exhibit button.)

Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer Area

Why was process 56 chosen as the deadlock victim?

- The deadlock priority for process 56 is the lowest.
- The execution priority for process 56 is the shortest.
- The database engine randomly chose process 56.
- The transaction rollback cost for process 56 is the last expensive.

Which lock type did process 56 attempt to acquire?

- a shared lock
- an exclusive lock
- a schema stability lock

Answer:

Explanation: References: <https://msdn.microsoft.com/en-us/library/ms186736.aspx>

NEW QUESTION 7

You have two database tables. Table1 is a partitioned table and Table 2 is a non-partitioned table.

Users report that queries take a long time to complete. You monitor queries by using Microsoft SQL Server Profiler. You observe lock escalation for Table1 and Table 2.

You need to allow escalation of Table1 locks to the partition level and prevent all lock escalation for Table2. Which Transact-SQL statement should you run for each table? To answer, drag the appropriate Transact-SQL statements to the correct tables. Each command 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.

Transact-SQL statements

- DBCC TRACEON(1211, -1)
- DBCC TRACEON(1224, -1)
- ALTER TABLE Table1 SET (LOCK_ESCALATION - DISABLE)
- ALTER TABLE Table2 SET (LOCK_ESCALATION - DISABLE)
- ALTER TABLE Table1 SET (LOCK_ESCALATION - AUTO)
- ALTER TABLE Table2 SET (LOCK_ESCALATION - AUTO)
- ALTER TABLE Table1 SET (LOCK_ESCALATION - TABLE)
- ALTER TABLE Table2 SET (LOCK_ESCALATION - TABLE)

Answer Area

Table	Transaction
Table1	Transact-SQL statement
Table2	Transact-SQL statement

Answer:

Explanation: Since SQL Server 2008 you can also control how SQL Server performs the Lock Escalation – through the ALTER TABLE statement and the property LOCK_ESCALATION. There are 3 different options available:

Box 1: Table1, Auto

The default option is TABLE, means that SQL Server *always* performs the Lock Escalation to the table level –even when the table is partitioned. If you have your table partitioned, and you want to have a Partition Level Lock Escalation (because you have tested your data access pattern, and you don't cause deadlocks with it), then you can change the option to AUTO. AUTO means that the Lock Escalation is performed to the partition level, if the table is partitioned, and otherwise to the table level.

Box 2: Table 2, DISABLE

With the option DISABLE you can completely disable the Lock Escalation for that specific table.

For partitioned tables, use the LOCK_ESCALATION option of ALTER TABLE to escalate locks to the HoBT level instead of the table or to disable lock escalation.

References:

<http://www.sqlpassion.at/archive/2014/02/25/lock-escalations/>

NEW QUESTION 8

You have a database named Sales.

You need to create a table named Customer that includes the columns described in the following table:

Column Name	Description	Masking requirement
CustomerID	An integer primary key for the customer record	None
FirstName	A string value that stored the first name of the customer	None
LastName	A string value that stored the last name of the customer	None
CreditLimit	A monetary value that stored the customer's credit limit	The entire field must be masked with a value between 1 and 100.
MobileNo	A ten-digit string value that stored the mobile number of the customer	The entire field must be masked with the exception of the first three digits.
Email	A string value that stored the email address of the customer	The entire field must be masked except the first letter of the user name, the @ symbol, and the top-level domain.

How should you complete the Transact SQL statement? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

```
CREATE TABLE Customer
```

```
(
```

```
    CustomerID int IDENTITY PRIMARY KEY,  
    FirstName varchar(100),  
    LastName varchar(100) NOT NULL,  
    CreditLimit money
```

▼
MASKED WITH (FUNCTION = 'partial(1, "000", 100)')
MASKED WITH (FUNCTION = 'default()')
MASKED WITH (FUNCTION = 'random(1,100)')

```
    MobileNO varchar(10)
```

▼
MASKED WITH (FUNCTION = 'partial(3, "XXXXXXX", 0)')
MASKED WITH (FUNCTION = 'default()')
MASKED WITH (FUNCTION = 'random(1,3)')

```
NULL,
```

```
    Email varchar(100)
```

▼
MASKED WITH (FUNCTION = 'email()')
MASKED WITH (FUNCTION = 'default()')
MASKED WITH (FUNCTION = 'partial(1, "XXXXXXX", 4)')
MASKED WITH (FUNCTION = 'partial(1, "@", 4)')

```
NULL,
```

```
);
```

Answer:

Explanation: Box 1: MASKED WITH (FUNCTION = 'default()')

The Default masking method provides full masking according to the data types of the designated fields. Example column definition syntax: Phone# varchar(12) MASKED WITH (FUNCTION = 'default()') NULL

Box 2: MASKED WITH (FUNCTION = 'partial(3, "XXXXXX", 0)')

The Custom String Masking method exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix examples:

PhoneNumber varchar(10) MASKED WITH (FUNCTION = 'partial(5, "XXXXXX", 0)')

Box 3: MASKED WITH (FUNCTION = 'email()')

The Email masking method which exposes the first letter of an email address and the constant suffix ".com", in the form of an email address. .aXXX@XXXX.com.

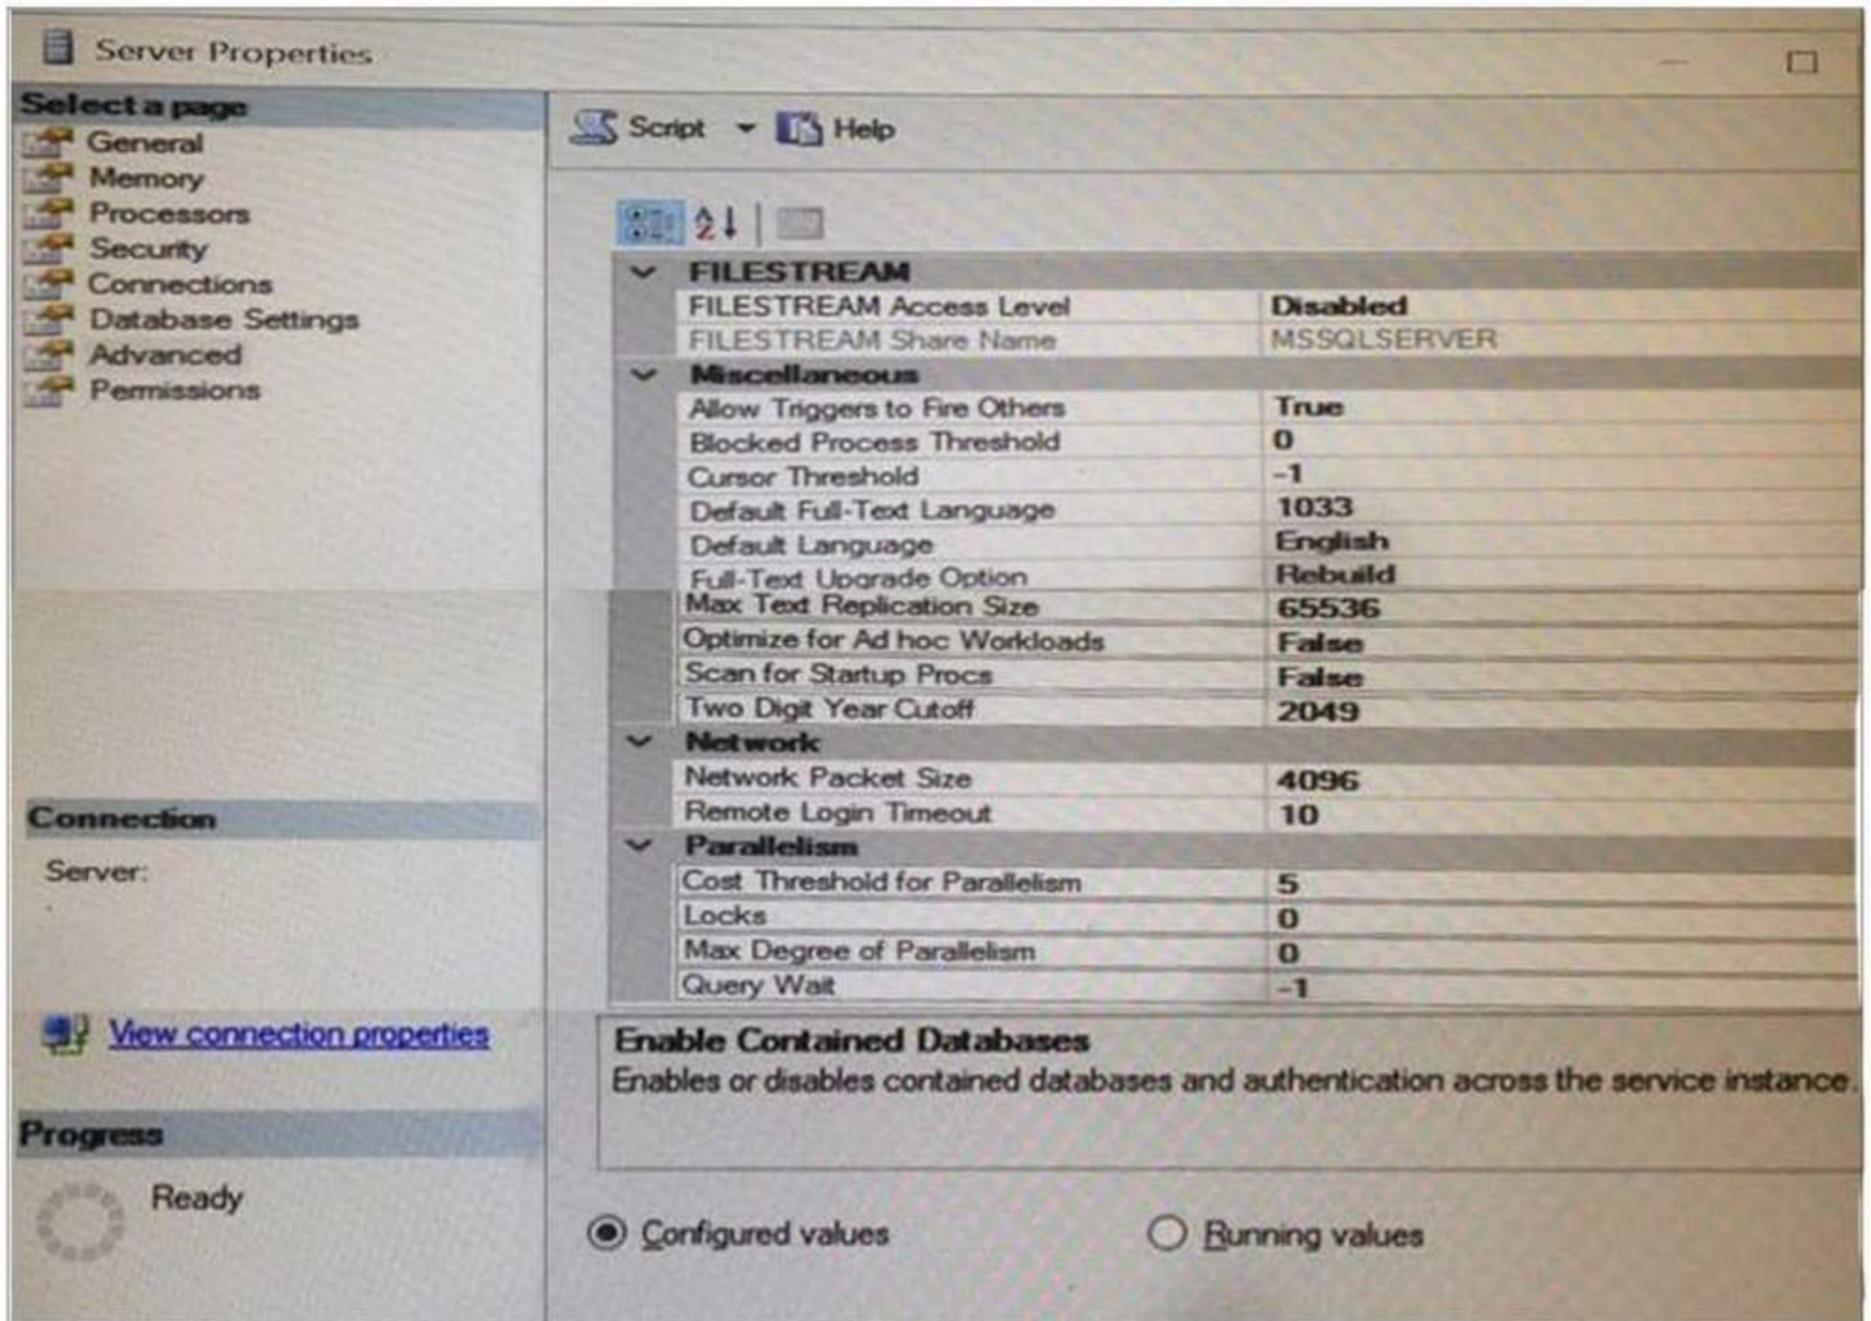
Example definition syntax: Email varchar(100) MASKEDWITH (FUNCTION = 'email()') NULL References: <https://msdn.microsoft.com/en-us/library/mt130841.aspx>

NEW QUESTION 9

You are reviewing the execution plans in the query plan cache. You observe the following:

- There are a large number of single use plans.
- There are a large number of simple execution plans that use multiple CPU cores. You need to configure the server to optimize query plan execution.

Which two setting should you modify on the properties page for the Microsoft SQL Server instance? To answer, select the appropriate settings in the answer area.



Answer:

Explanation: * Optimize for ad hoc workloads

The optimize for ad hoc workloads option is used to improve the efficiency of the plan cache for workloads that contain many single use ad hoc batches. When this option is set to 1, the Database Engine stores a small compiled plan stub in the plan cache when a batch is compiled for the first time, instead of the full compiled plan. This helps to relieve memory pressure by not allowing the plan cache to become filled with compiled plans that are not reused.

* Cost Threshold for Parallelism

Use the cost threshold for parallelism option to specify the threshold at which Microsoft SQL Server creates and runs parallel plans for queries. SQL Server creates and runs a parallel plan for a query only when the estimated cost to run a serial plan for the same query is higher than the value set in cost threshold for parallelism. The cost refers to an estimated elapsed time in seconds required to run the serial plan on a specific hardware configuration.

5 means 5 seconds, but is is 5 seconds on a machine internal to Microsoft from some time in the 1990s. There's no way to relate it to execution time on your current machine, so we treat it as a pure number now. Raising it to 50 is a common suggestion nowadays, so that more of your simpler queries run on a single thread.

NEW QUESTION 10

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a database that contains a table named Employees. The table stores information about the employees of your company.

You need to implement and enforce the following business rules:

- Limit the values that are accepted by the Salary column.
- Prevent salaries less than \$15,000 and greater than \$300,000 from being entered.
- Determine valid values by using logical expressions.
- Do not validate data integrity when running DELETE statements. Solution: You implement a FOR UPDATE trigger on the table. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: References:

<http://stackoverflow.com/questions/16081582/difference-between-for-update-of-and-for-update>

NEW QUESTION 10

Note: This question is part of a series of questions that use the same or similar answer choices. An Answer choice may be correct for more than one question in the series. Each question independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run during off-peak hours. You observe that many deadlocks appear to be happening during specific times of the day. You need to monitor the SQL environment and capture the information about the processes that are causing the deadlocks. What should you do?

- A. Create a sys.dm_os_waiting_tasks query.
- B. Create a sys.dm_exec_sessions query.
- C. Create a PerformanceMonitor Data Collector Set.
- D. Create a sys.dm_os_memory_objects query.
- E. Create a sp_configure 'max server memory' query.
- F. Create a SQL Profiler trace.
- G. Create a sys.dm_os_wait_stats query.
- H. Create an Extended Event.

Answer: F

Explanation: To view deadlock information, the Database Engine provides monitoring tools in the form of two trace flags, and the deadlock graph event in SQL Server Profiler.

Trace Flag 1204 and Trace Flag 1222

When deadlocks occur, trace flag 1204 and trace flag 1222 return information that is captured in the SQL Server error log. Trace flag 1204 reports deadlock information formatted by each node involved in the deadlock. Trace flag 1222 formats deadlock information, first by processes and then by resources. It is possible to enable both trace flags to obtain two representations of the same deadlock event.

References: [https://technet.microsoft.com/en-us/library/ms178104\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms178104(v=sql.105).aspx)

NEW QUESTION 11

You have an existing Microsoft SQL Trace script. You plan to convert the script to an Extended Events session. You need to collect the trace ID and other required information. Which system table should you use?

- A. dbo.sysessions
- B. trace_xenction_map
- C. db
- D. syinotifications
- E. sysdbmaintplan.

Answer: A

NEW QUESTION 14

You have a database that users query frequently.

The users report that during peak business hours, the queries take longer than expected to execute.

A junior database administrator uses Microsoft SQL Server Profiler on the database server to trace the session activities.

While performing the trace, the performance of the database server worsens, and the server crashes.

You need to recommend a solution to collect the query run times. The solution must minimize the impact on the resources of the database server.

What should you recommend?

- A. Increase the free space on the system drive of the database server, and then use SQL Server Profiler on the server to trace the session activities.
- B. Collect session activity data by using SQL Server Extended Events.
- C. Clean up tempdb, and then use SQL Server Profiler on the datafile server to trace the session activities.
- D. Collect performance data by using a Data Collector Set (DCS) in Performance Monitor.

Answer: A

NEW QUESTION 16

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a database named DB1 that includes a table named sales.Orders. You grant a user named User1 select permissions on the sales schema.

You need to ensure that User1 can select data from the sales.orders table without specifying the schema name in any Transact SQL statements.

Solution: You move the sales.orders table to the dbo schema. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 21

You have the following stored procedure that is called by other stored procedures and applications:

```
CREATE PROCEDURE UpdateCustomer @CustomerId INT
AS
BEGIN
    EXEC ProcessCustomer_Internal @CustomerId
    DECLARE @Status INT
    SELECT @Status = Status FROM Customer WHERE CustomerId =
@CustomerId
END
```

You need to modify the stored procedure to meet the following requirements:

Always return a value to the caller.

Return 0 if @Status is NULL.

Callers must be able to use @Status as a variable.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Replace NULL values with 0. Add a PRINT statement to return @Status.
- B. Add a RETURN statement.
- C. Replace NULL values with 0. Add an output parameter to return @Status.
- D. Replace NULL values with 0. Add a SELECT statement to return @Status.
- E. Add a PRINT statement.
- F. Add a SELECT statement to return @Status.
- G. Add an output parameter to return @Status.

Answer: BC

Explanation: There are three ways of returning data from a procedure to a calling program: result sets, output parameters, and return codes.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/stored-procedures/return-data-from-a-stored-procedure>

NEW QUESTION 23

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

You have a database named DB1 that contains the following tables: Customer, CustomerToAccountBridge, and CustomerDetails. The three tables are part of the Sales schema. The database also contains a schema named Website. You create the Customer table by running the following Transact-SQL statement:

```
CREATE TABLE Customer
(
    CustomerNumber int NOT NULL,
    CustomerName varchar(50) NOT NULL,
    CreateDate date NOT NULL,
    Gender bit,
    Address varchar(50)
    City varchar(50)
    State char(2),
    CustomerStatus bit NOT NULL,
    MaritalStatus bit,
    Segment varchar(5),
    CountryCode char(2),
    Birthday date,
    PostalCode char(5),
    PhoneNumber varchar(20),
    Account1 char(7),
    Account1Status bit,
    Account2 char(7),
    Account2Status bit,
    CONSTRAINT PK_Customer PRIMARY KEY CLUSTERED (CustomerNumber)
);
```

The value of the CustomerStatus column is equal to one for active customers. The value of the Account1Status and Account2Status columns are equal to one for active accounts. The following table displays selected columns and rows from the Customer table.

Customer ID	CustomerName	Gender	Account1	Account1Status	Account2	Account2Status
101	Name A	0	0001001	0	0001002	1
102	Name B	1	0002001	1	0002002	0
103	Name C	0	0003001	1	0003002	1

You plan to create a view named Website.Customer and a view named Sales.FemaleCustomers. Website.Customer must meet the following requirements:

- * Allow users access to the CustomerName and CustomerNumber columns for active customers.
- * Allow changes to the columns that the view references. Modified data must be visible through the view.
- * Prevent the view from being published as part of Microsoft SQL Server replication. Sales.FemaleCustomers must meet the following requirements:
- * Allow users access to the CustomerName, Address, City, State and PostalCode columns.
- * Prevent changes to the columns that the view references.
- * Only allow updates through the views that adhere to the view filter.

You have the following stored procedures: spDeleteCustAcctRelationship and spUpdateCustomerSummary. The spUpdateCustomerSummary stored procedure was created by running the following Transact-SQL statement:

```
CREATE PROCEDURE uspUpdateCustomerSummary
@CustomerId INT
AS
BEGIN
    SET NOCOUNT on;
    UPDATE CustomerDetails SET TotalDepositAccountCount = TotalDepositAccountCount + 1 WHERE CustomerID = @CustomerId;
    BEGIN TRAN;
    BEGIN TRY
        UPDATE CustomerDetails SET TotalAccountCount = TotalAccountCount + 1 WHERE CustomerID = @CustomerId;
    END TRY
    BEGIN CATCH
        IF @@TRANCOUNT > 0
            ROLLBACK TRAN;
    END CATCH
    IF @@TRANCOUNT > 0
        COMMIT TRAN;
END
```

You run the spUpdateCustomerSummary stored procedure to make changes to customer account summaries. Other stored procedures call the spDeleteCustAcctRelationship to delete records from the CustomerToAccountBridge table. When a procedure calls spDeleteCustAcctRelationship, if the calling stored procedure has already started an active transaction, all the deletions made by the spDeleteCustAcctRelationship stored procedure must be committed by the caller; otherwise changes must be committed within the spDeleteCustAcctRelationship stored procedure. If any error occurs during the delete operation, only the deletes made by the spDeleteCustAcctRelationship stored procedure must be rolled back and the status must be updated. You need to complete the stored procedure to ensure all the requirements are met. How should you complete the procedure? To answer, drag the Transact-SQL segments to the correct location. Each transact-SQL segment 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. NOTE: Each correct selection is worth one point.

Transact-SQL segments

- BEGIN TRANSACTION;
- COMMIT TRANSACTION;
- ROLLBACK TRANSACTION
- SAVE TRANSACTION SavePoint;
- ROLLBACK TRANSACTION SavePoint;

Answer Area

```
CREATE PROCEDURE spDeleteCustAcctRelationship
(@CustomerId int)
AS
BEGIN
    SET NOCOUNT ON;
    DECLARE @TranStartedAlready CHAR(1) = "N";
    DECLARE @status INT = 0;

    IF @@TRANCOUNT > 0
        SET @TranStartedAlready = "Y";
    IF @TranStartedAlready = "Y"
        Transact-SQLstatement
    ELSE
        Transact-SQLstatement

    BEGIN TRY
        DELETE FROM CustomerToAccountBridge WHERE CustomerId = @CustomerId;
        IF @TranStartedAlready = "N"
            BEGIN
                Transact-SQLstatement
                SET @Status = 0;
            END
    END TRY
    BEGIN CATCH
        IF @TranStartedAlready = "N"
            Transact-SQLstatement
        ELSE
            IF XACT_STATE() <> -1
                Transact-SQLstatement
            SET @Status = 1;
    END CATCH;
END;
```

Answer:

Explanation: Savepoints offer a mechanism to roll back portions of transactions. You create a savepoint using the SAVE TRANSACTION savepoint_name statement. Later, you execute a ROLLBACK TRANSACTION savepoint_name statement to roll back to the savepoint instead of rolling back to the start of the transaction. References: [https://technet.microsoft.com/en-us/library/ms178157\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms178157(v=sql.105).aspx)

NEW QUESTION 24

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals. You have a table that has a clustered index and a nonclustered index. The indexes use different columns from the table. You have a query named Query1 that uses the nonclustered index. Users report that Query1 takes a long time to report results. You run Query1 and review the following statistics for an index seek operation:

Index Seek (NonClustered)

Scan a particular range of rows from a nonclustered index.

Physical Operation	Index Seek
Logical Operation	Index Seek
Actual Execution Mode	Row
Actual Number of Rows	3571454
Actual Number of Batches	0
Estimated I/O Cost	0.0093577
Estimated Operator Cost	0.0107304 (0%)
Estimated CPU Cost	0.0013727
Estimated Subtree Cost	0.0107304
Number of Executions	8
Estimated Number of Rows	0
Estimated Row Size	19 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	100

You need to resolve the performance issue. Solution: You drop the nonclustered index.
 Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 27

You are creating the following two stored procedures:

- A. a natively-compiled stored procedure
 - An interpreted stored procedure that accesses both disk-based and memory-optimized tables
- Both stored procedures run within transactions.

You need to ensure that cross-container transactions are possible. Which setting or option should you use?

- A. the SET TRANSACTION_READ_COMMITTED isolation level for the connection
- B. the SERIALIZABLE table hint on disk-based tables
- C. the SET MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT=ON option for the database
- D. the SET MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT=OFF option for the database

Answer: C

Explanation: Provide a supported isolation level for the memory-optimized table using a table hint, such as WITH (SNAPSHOT). The need for the WITH (SNAPSHOT) hint can be avoided through the use of the database option MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT. When this option is set to ON, access to a memory-optimized table under a lower isolation level is automatically elevated to SNAPSHOT isolation.

NEW QUESTION 30

You have a trigger named CheckTriggerCreation that runs when a user attempts to create a trigger. The CheckTriggerCreation trigger was created with the ENCRYPTION option and additional proprietary business logic.

You need to prevent users from running the ALTER and DROP statements or the sp_tableoption stored procedure.

Which three Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

Commands

```
DISABLE TRIGGER CheckTriggerCreation ON DATABASE;
```

```
ENABLE TRIGGER CheckTriggerCreation ON DATABASE;
```

```
CREATE TRIGGER CheckTriggerCreation ON DATABASE FOR CREATE_TRIGGER AS RAISERROR ('Error message', 10, 1) ROLLBACK GO
```

```
CREATE TRIGGER CheckTableChanges ON DATABASE FOR ALTER_TABLE, DROP_TABLE AS RAISERROR ('Error message', 10, 1) ROLLBACK GO
```

```
DROP TRIGGER CheckTrigerCreation ON DATABASE;
```

```
CREATE TRIGGER CheckTableChanges ON DATABASE FOR DDL_TABLE_EVENTS AS RAISERROR ('Error message', 10, 1) ROLLBACK GO
```

Answer Area



Answer:

Explanation:

Commands

```
DISABLE TRIGGER CheckTriggerCreation ON DATABASE;
```

```
ENABLE TRIGGER CheckTriggerCreation ON DATABASE;
```

```
CREATE TRIGGER CheckTriggerCreation ON DATABASE FOR CREATE_TRIGGER AS RAISERROR ('Error message', 10, 1) ROLLBACK GO
```

```
CREATE TRIGGER CheckTableChanges ON DATABASE FOR ALTER_TABLE, DROP_TABLE AS RAISERROR ('Error message', 10, 1) ROLLBACK GO
```

```
DROP TRIGGER CheckTrigerCreation ON DATABASE;
```

```
CREATE TRIGGER CheckTableChanges ON DATABASE FOR DDL_TABLE_EVENTS AS RAISERROR ('Error message', 10, 1) ROLLBACK GO
```

Answer Area

```
DISABLE TRIGGER CheckTriggerCreation ON DATABASE;
```

```
CREATE TRIGGER CheckTableChanges ON DATABASE FOR DDL_TABLE_EVENTS AS RAISERROR ('Error message', 10, 1) ROLLBACK GO
```

```
ENABLE TRIGGER CheckTriggerCreation ON DATABASE;
```



NEW QUESTION 32

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
 You have a database named DB1 that contains the following tables: Customer, CustomerToAccountBridge, and CustomerDetails. The three tables are part of the Sales schema. The database also contains a schema named Website. You create the Customer table by running the following Transact-SQL statement:

```
CREATE TABLE Customer
(
    CustomerNumber int NOT NULL,
    CustomerName varchar(50) NOT NULL,
    CreateDate date NOT NULL,
    Gender bit,
    Address varchar(50)
    City varchar(50)
    State char(2),
    CustomerStatus bit NOT NULL,
    MaritalStatus bit,
    Segment varchar(5),
    CountryCode char(2),
    Birthday date,
    PostalCode char(5),
    PhoneNumber varchar(20),
    Account1 char(7),
    Account1Status bit,
    Account2 char(7),
    Account2Status bit,
    CONSTRAINT PK_Customer PRIMARY KEY CLUSTERED (CustomerNumber)
);
```

The value of the CustomerStatus column is equal to one for active customers. The value of the Account1Status and Account2Status columns are equal to one for active accounts. The following table displays selected columns and rows from the Customer table.

Customer ID	CustomerName	Gender	Account1	Account1Status	Account2	Account2Status
101	Name A	0	0001001	0	0001002	1
102	Name B	1	0002001	1	0002002	0
103	Name C	0	0003001	1	0003002	1

You plan to create a view named Website.Customer and a view named Sales.FemaleCustomers. Website.Customer must meet the following requirements:

- * Allow users access to the CustomerName and CustomerNumber columns for active customers.
- * Allow changes to the columns that the view references. Modified data must be visible through the view.
- * Prevent the view from being published as part of Microsoft SQL Server replication. Sales.FemaleCustomers must meet the following requirements:
- * Allow users access to the CustomerName, Address, City, State and PostalCode columns.
- * Prevent changes to the columns that the view references.
- * Only allow updates through the views that adhere to the view filter.

You have the following stored procedures: spDeleteCustAcctRelationship and spUpdateCustomerSummary. The spUpdateCustomerSummary stored procedure was created by running the following Transact-SQL statement:

```
CREATE PROCEDURE uspUpdateCustomerSummary
@CustomerId INT
AS
BEGIN
    SET NOCOUNT on;
    UPDATE CustomerDetails SET TotalDepositAccountCount = TotalDepositAccountCount + 1 WHERE CustomerID = @CustomerId;
    BEGIN TRAN;
    BEGIN TRY
        UPDATE CustomerDetails SET TotalAccountCount = TotalAccountCount + 1 WHERE CustomerID = @CustomerId;
    END TRY
    BEGIN CATCH
        IF @@TRANCOUNT > 0
            ROLLBACK TRAN;
    END CATCH
    IF @@TRANCOUNT > 0
        COMMIT TRAN;
```

You run the spUpdateCustomerSummary stored procedure to make changes to customer account summaries. Other stored procedures call the spDeleteCustAcctRelationship to delete records from the CustomerToAccountBridge table. You need to create Website Customer.

How should you complete the view definition? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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.

NOTE: Each correct selection is worth one point. NOTE: Each correct selection is worth one point.

Transact-SQL segments

- WITH SCHEMABINDING
- WITH ENCRYPTION
- WITH CHECK OPTION
- WITH VIEW_METADATA

Answer area

```
CREATE VIEW Website.Customer
    Transact-SQL statement
AS SELECT s.CustomerID, s.CustomerName
FROM Sales.Customer as s
WHERE s.CustomerStatus=1
    Transact-SQL statement
```

Answer:

Explanation: Box 1: WITH ENCRYPTION

Using WITH ENCRYPTION prevents the view from being published as part of SQL Server replication. Box 2: WITH CHECK OPTION
 CHECK OPTION forces all data modification statements executed against the view to follow the criteria set within select_statement. When a row is modified through a view, the WITH CHECK OPTION makes sure the data remains visible through the view after the modification is committed.

Note: Website.Customer must meet the following requirements:

NEW QUESTION 37

Note: This question is part of a series of questions that use the same answer choices. An answer choice may be correct for more than one question on the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You work on an OLTP database that has no memory-optimized file group defined.

You have a table named tblTransaction that is persisted on disk and contains the information described in the following table:

Item	Name	Data Type	Nullable	Notes
Column	TransactionDate	Date	No	For each transaction date, there are only about 100,000 records. The table contains over one billion records in total.
Column	SequenceNo	bigint	No	Uniquely identifies a transaction record within a date
Column	AccountId	int	No	
Column	ValueType	char(3)	No	
Column	Amount	decimal(20,2)	Yes	
	IX_ValueType			Nonclustered columnstore index on the ValueType column.

Users report that the following query takes a long time to complete.

```
SELECT TransactionDate, COUNT(*) AS TotalCount FROM tblTransaction
WHERE TransactionDate - DATEADD(D, -1, CONVERT (DATE, CONVERT (VARCHAR (8),
GETDATE ( ), 112) 112))
GROUP BY TransactionDate;
```

You need to create an index that:

- improves the query performance
- does not impact the existing index
- minimizes storage size of the table (inclusive of index pages). What should you do?

- A. Create a clustered index on the table.
- B. Create a nonclustered index on the table.
- C. Create a nonclustered filtered index on the table.
- D. Create a clustered columnstore index on the table.
- E. Create a nonclustered columnstore index on the table.
- F. Create a hash index on the table.

Answer: C

Explanation: A filtered index is an optimized nonclustered index, especially suited to cover queries that select from a well-defined subset of data. It uses a filter predicate to index a portion of rows in the table. A well-designed filtered index can improve query performance, reduce index maintenance costs, and reduce index storage costs compared with full-table indexes.

NEW QUESTION 40

Note: This question is part of a series of questions that present the same scenario. Each question in this series contains a unique solution. Determine whether the solution meets the stated goals.

Your company has employees in different regions around the world.

You need to create a database table that stores the following employee attendance information:

- Employee ID
- date and time employee checked in to work
- date and time employee checked out of work

Date and time information must be time zone aware and must not store fractional seconds. Solution: You run the following Transact-SQL statement:

```
CREATE TABLE [dbo].[EmployeeAttendance] (  
    EmployeeID int NOT NULL,  
    DateChekedIn datetimeoffset NOT NULL,  
    DateCheclOut datetimeoffset NOT NULL)
```

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: Datetimeoffset, not datetimeoffset, defines a date that is combined with a time of a day that has time zone awareness and is based on a 24-hourclock.

Syntax: datetimeoffset [(fractional seconds precision)]

For the use "datetimeoffset", the Fractional seconds precision is 7. References: <https://msdn.microsoft.com/en-us/library/bb630289.aspx>

NEW QUESTION 41

You have several real-time applications that constantly update data in a database. The applications run more than 400 transactions per second that insert and update new metrics from sensors.

A new web dashboard is released to present the data from the sensors. Engineers report that the applications take longer than expected to commit updates. You need to change the dashboard queries to improve concurrency and to support reading uncommitted data. What should you do?

- A. Use the NOLOCK option.
- B. Execute the DBCC UPDATEUSAGE statement.
- C. Use the max worker threads Option.
- D. Use a table-valued parameter.
- E. Set SET ALLOW SNAPSHOT ISOLATION to ON.
- F. Set SET XACTJVBORT to ON.
- G. Execute the alter table ti set (Lock_ESCALATION = auto); statement.
- H. Use the output parameters.

Answer: C

NEW QUESTION 43

You are developing an ETL process to cleanse and consolidate incoming data. The ETL process will use a reference table to identify which data must be cleansed in the target table. The server that hosts the tables restarts daily.

You need to minimize the amount of time it takes to execute the query and the amount of time it takes to populate the reference table.

What should you do?

- A. Convert the target table to a memory-optimized tabl
- B. Create a natively compiled stored procedure to cleanse and consolidate the data.
- C. Convert the reference table to a memory-optimized tabl
- D. Set the durability option toSCHEMA_AND_DATA>
- E. Create a native compiled stored procedure to implement the ETL process for both tables.
- F. Convert the reference table to a memory-optimized tabl
- G. Set the durability option to SCHEMA_ONLY.

Answer: D

NEW QUESTION 45

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a database named DB1 that includes a table named sales .orders. You grant a user named User1 select permissions on the Sales schema.

You need to ensure that User1 can select data from the sales .orders table without specifying the schema name in any Transact -SQL statements.

Solution: You set the default database for User1 to DB1. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 46

You use Microsoft SQL Server Profile to evaluate a query named Query1. The Profiler report indicates the following issues:

At each level of the query plan, a low total number of rows are processed.

The query uses many operations. This results in a high overall cost for the query. You need to identify the information that will be useful for the optimizer. What should you do?

- A. Start a SQL Server Profiler trace for the event class Performance statistics in the Performance eventcategory.
- B. Create one Extended Events session with the sqlserver.missing_column_statistics event added.
- C. Start a SQL Server Profiler trace for the event class Soft Warnings in the Errors and Warnings event category.
- D. Create one Extended Events session with the sqlserver.error_reported event added.

Answer: A

Explanation: The Performance Statistics event class can be used to monitor the performance of queries, stored procedures, and triggers that are executing. Each of the six event subclasses indicates an event in the lifetime of queries, stored procedures, and triggers within the system. Using the combination of these event subclasses and the associated sys.dm_exec_query_stats, sys.dm_exec_procedure_stats and sys.dm_exec_trigger_stats dynamic management views, you can reconstitute the performance history of any given query, stored procedure, or trigger.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/event-classes/performance-statistics-event-class?view=>

NEW QUESTION 48

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a database that contains a table named Employees. The table stores information about the employees of your company.

You need to implement and enforce the following business rules:

- Limit the values that are accepted by the Salary column.
- Prevent salaries less than \$15,000 and greater than \$300,000 from being entered.
- Determine valid values by using logical expressions.
- Do not validate data integrity when running DELETE statements.

Solution: You implement cascading referential integrity constraints on the table. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation: References: [https://technet.microsoft.com/en-us/library/ms186973\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms186973(v=sql.105).aspx)

NEW QUESTION 51

You are planning a set of stored procedures that must be able to access memory-optimized tables. You need to optimize the performance of the stored procedures. Which statement should you include in the stored procedure definitions?

- A. WITH RECOMPILE
- B. WITH NATIVE_COMPILATION
- C. WITH EXECUTE AS SELF
- D. WITH NO INFOMSG5

Answer: D

NEW QUESTION 54

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

The Account table was created by using the following Transact-SQL statement:

```
CREATE TABLE Account
(
    AccountNumber int NOT NULL,
    ProductCode char(2) NOT NULL,
    Status tinyint NOT NULL,
    OpenDate date NOT NULL,
    CloseDate date,
    Balance decimal(15,2),
```

There are more than 1 billion records in the Account table. The AccountNumber column uniquely identifies each account. The productcode column has 100 different values. The values are evenly distributed in the table. Table statistics are refreshed and up to date.

You frequently run the following Transact-SQL select statements:

```
SELECT ProductCode, SUM(Balance) AS TotalSUH FROM Account WHERE ProductCode <> 'CD' GROUP BY ProductCode; SELECT AccountNumber, Balance FROM Account WHERE ProductCode = 'CD';
```

You must avoid table scans when you run the queries. You need to create one or more indexes for the table. Solution: You run the following Transact-SQL statements:

```
CREATE CLUSTERED INDEX PK_Account ON Account(AccountNumber);
CREATE NONCLUSTERED INDEX IX_Account_ProductCode ON Account(ProductCode) INCLUDE (Balance);
```

Does the solution meet the goal?

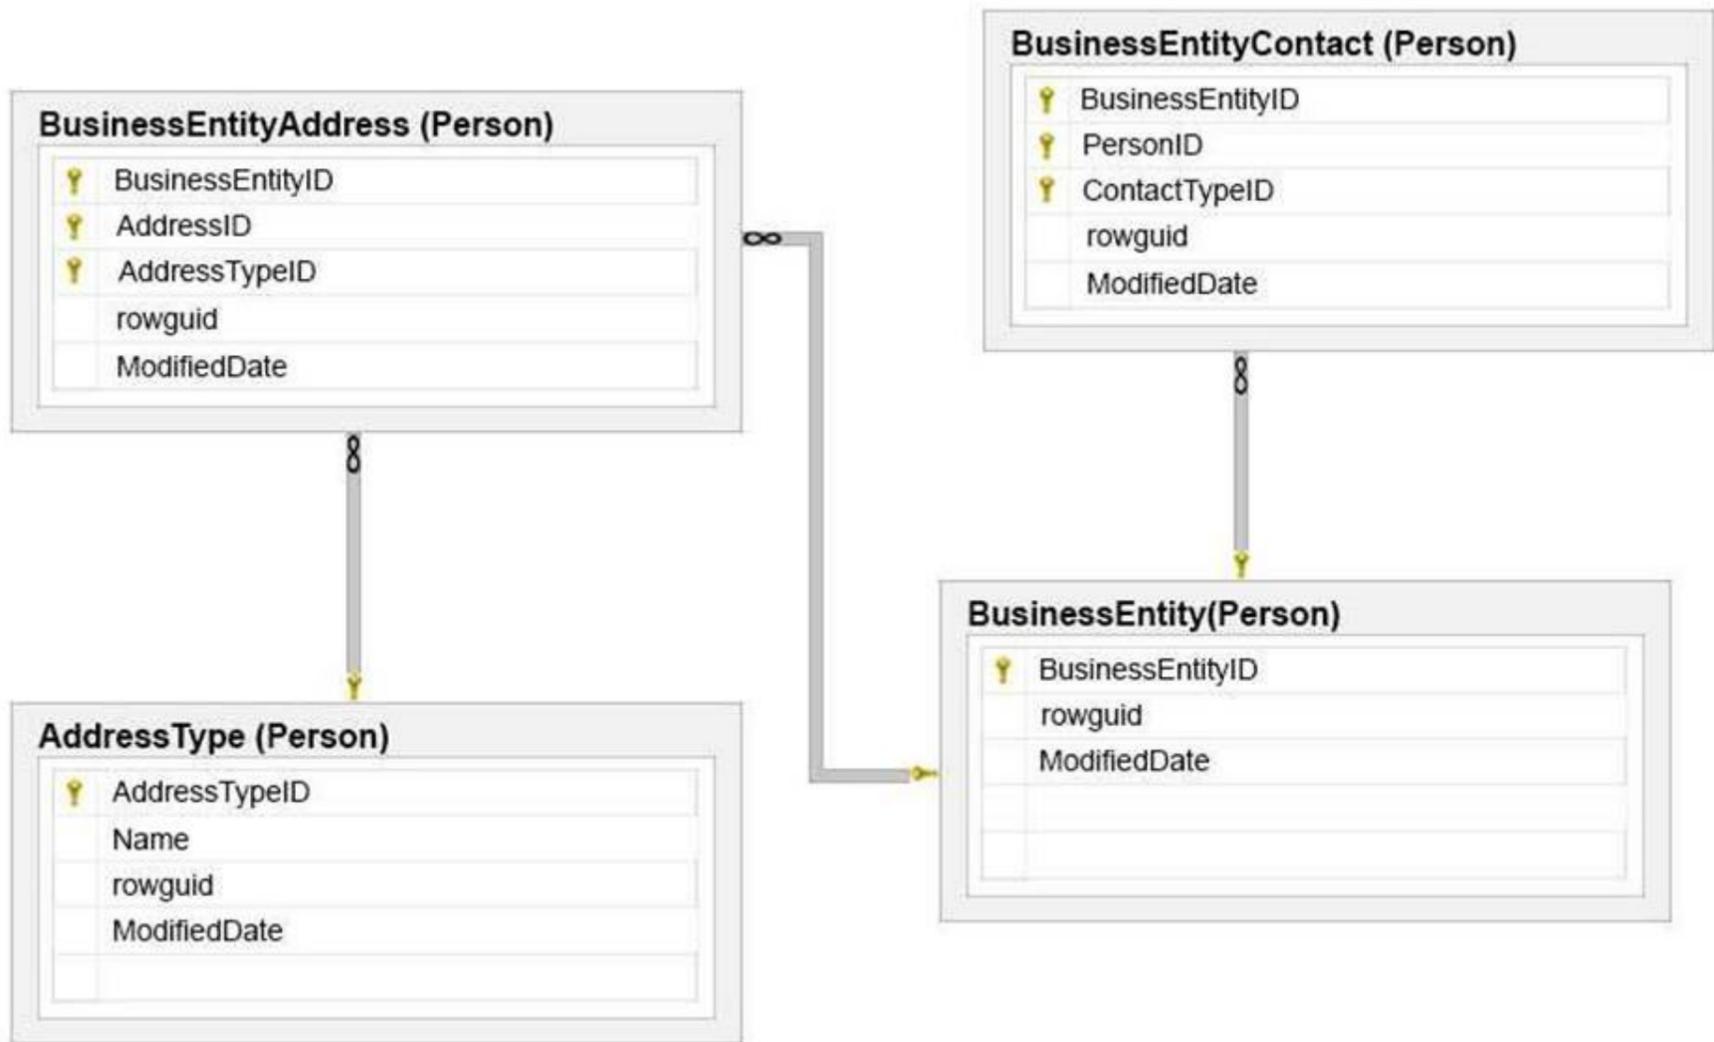
- A. Yes

B. No

Answer: A

NEW QUESTION 56

You are creating a stored procedure which will insert data into the table shown in the Database schema exhibit. (Click the exhibit button.)



You need to insert a new customer record into the tables as a single unit of work.

In which order should you use the Transact-SQL segments to develop the solution? To answer, move the appropriate Transact-SQL segments 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 of the correct orders you select.

Transact-SQL segments

- COMMIT TRANSACTION
- INSERT INTO Person.AddressType
- INSERT INFO Person.BusinessEntityAddress
- INSERT INTO Person. BusinessEntity
- BEGIN TRANSACTION
- INSERT INTO Person.BusinessEntityContact

Answer Area



Answer:

Explanation: The entities on the many side, of the 1-many relations, must be added before we add the entities on the 1-side. We must insert new rows into BusinessEntityContact and BusinessEntityAddress tables, before we insert the corresponding rows into the BusinessEntity and AddressType tables.

NEW QUESTION 57

You run the following Transact-SQL statements:

```
CREATE TABLE Customers(
  CustomerID INT NOT NULL IDENTITY PRIMARY KEY CLUSTERED,
  CustomerName NVARCHAR (100) UNIQUE NOT NULL
)

CREATE TABLE Orders(
  OrderID INT NOT NULL IDENTITY PRIMARY KEY CLUSTERED,
  CustomerID INT NOT NULL REFERENCES Customers (CustomerID),
  OrderDate DATE NOT NULL
)
CREATE VIEW v_CustomerOrder
AS SELECT
  b.CustomerName, a.OrderID, a.OrderDate,
  (SELECT COUNT(*) FROM Orders c WHERE c.CustomerID = a.CustomerID) AS CustomerOrderCount
FROM Orders a
INNER JOIN Customers b ON a.CustomerID = b.CustomerID
```

Records must only be added to the orders table by using the view. If a customer name does not exist then a new customer name must be created. You need to ensure that you can insert rows into the orders table by using the view. What should you do?

- A. Add the CustomerID column from the Orders table and the WITH CHECK OPTION statement to the view.
- B. Create an INSTEAD of trigger on the view.
- C. Add the WITH SCHEMABINDING statement to the view statement and create a clustered index on the view.
- D. Remove the subquery from the view, add the WITH SCHEMABINDING statement, and add a trigger to the Orders table to perform the required logic.

Answer: C

NEW QUESTION 58

You have a database that contains both disk-based and memory-optimized tables. You need to create two modules. The modules must meet the requirements described in the following table.

Module	Requirements
Module 1	<ul style="list-style-type: none"> • must be encrypted by using the ENCRYPTION option • must support updates on both disk-based and memory-optimized tables • must support OUTPUT parameters
Module 2	<ul style="list-style-type: none"> • must access only memory-optimized tables • must support updates on memory-optimized tables • must support heavy aggregations with highest performance • must support OUTPUT parameters.

Which programming object should you use for each module? To answer, select the appropriate object types in the answer area.

Answer Area

Module	Object type
Module1	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #cccccc; height: 20px; margin-bottom: 5px;"></div> <div style="display: flex; justify-content: flex-end; align-items: center;">▼</div> <ul style="list-style-type: none"> interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function </div>
Module2	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #cccccc; height: 20px; margin-bottom: 5px;"></div> <div style="display: flex; justify-content: flex-end; align-items: center;">▼</div> <ul style="list-style-type: none"> interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function </div>

Answer:

Explanation: Module 1: Interpreted stored procedure

An interpreted stored procedure can access both disk-based and memory-optimized tables. Module 2: Natively compiled stored procedure

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.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/natively-compiled-stored-procedures?>

NEW QUESTION 59

You create a database table named FactSales by running the following Transact-SQL statements:

```
CREATE TABLE FactSales (
    DateKey [int] NOT NULL,
    ProductKey [int] NOT NULL,
    CustomerKey [int] NOT NULL,
    EmployeeKey [int] NOT NULL,
    SalesAmount money,
    SalesCost money
    INDEX ix_FactSales CLUSTERED (DateKey));

CREATE INDEX ix_FactSales_ProductKey ON FactSales (ProductKey);
CREATE INDEX ix_FactSales_CustomerKey ON FactSales (CustomerKey);
CREATE INDEX ix_FactSales_EmployeeKey ON FactSales (EmployeeKey);
```

You must optimize the indexes without making changes to the ix_FactSales_EmployeeKey index. You need to implement a columnstore index for the table. How should you complete the Transact-SQL code? To answer, select the appropriate Transact-SQL segments in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```

DROP INDEX ix_FactSales_ProductKey ON FactSales
DROP INDEX ix_FactSales_EmployeeKey ON FactSales
DROP INDEX ix_FactSales ON FactSales

CREATE CLUSTERED COLUMNSTORE INDEX ix FactSales ON FactSales
DROP INDEX ix_FactSales_CustomerKey ON FactSales
DROP INDEX ix_FactSales_EmployeeKey ON FactSales
DROP INDEX ix_FactSales ON FactSales

CREATE CLUSTERED COLUMNSTORE INDEX ix FactSales ON FactSales
CREATE CLUSTERED COLUMNSTORE INDEX ix_FactSales ON FactSales WITH (DROP_EXISTING = ON)
ALTER INDEX ix_FactSales_EmployeeKey ON FactSales REBUILD
CREATE INDEX ix_FactSales_EmployeeKey ON FactSales (EmployeeKey);
    
```

Answer:

Explanation:

Answer Area

```

DROP INDEX ix_FactSales_ProductKey ON FactSales
DROP INDEX ix_FactSales_EmployeeKey ON FactSales
DROP INDEX ix_FactSales ON FactSales

CREATE CLUSTERED COLUMNSTORE INDEX ix FactSales ON FactSales
DROP INDEX ix_FactSales_CustomerKey ON FactSales
DROP INDEX ix_FactSales_EmployeeKey ON FactSales
DROP INDEX ix_FactSales ON FactSales

CREATE CLUSTERED COLUMNSTORE INDEX ix FactSales ON FactSales
CREATE CLUSTERED COLUMNSTORE INDEX ix FactSales ON FactSales WITH (DROP_EXISTING = ON)
ALTER INDEX ix_FactSales_EmployeeKey ON FactSales REBUILD
CREATE INDEX ix_FactSales_EmployeeKey ON FactSales (EmployeeKey);
    
```

NEW QUESTION 62

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a database that contains a table named Employees. The table stores information about the employees of your company.

You need to implement the following auditing rules for the Employees table:

- Record any changes that are made to the data in the Employees table.
- Customize the data recorded by the audit operations.

Solution: You implement a user-defined function on the Employees table. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation: SQL Server 2016 provides two features that track changes to data in a database: change data capture and change tracking. These features enable applications to determine the DML changes (insert, update, and delete operations) that were made to user tables in a database.

Change data is made available to change data capture consumers through table-valued functions (TVFs). References: <https://msdn.microsoft.com/en-us/library/cc645858.aspx>

NEW QUESTION 64

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questions sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have a database that is 130 GB and contains 500 million rows of data. Granular transactions and mass batch data imports change the database frequently throughout the day. Microsoft SQL Server Reporting Services (SSRS) uses the database to generate various reports by using several filters. You discover that some reports time out before they complete. You need to reduce the likelihood that the reports will time out. Solution: You partition the largest tables. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 67

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

Table name	Description
TBL1	<ul style="list-style-type: none"> • The table has 25 columns. • The table will contain 10 million records. • Approximately 100,000 records will be inserted monthly.
TBL2	<ul style="list-style-type: none"> • The table has 25 columns. • The table will contain 100,000 records. • The frequency of inserting, updating, and deleting records is low.

You have a Microsoft SQL Server database named DB1 that contains the following tables: You frequently run the following queries:

```
SELECT *
FROM TBL1
WHERE Column1 BETWEEN '01/01/2016' AND '30/04/2016'

SELECT Column5, Column6
FROM TBL2
WHERE Column2 = 'ABC156XYZ'
```

There are no foreign key relationships between TBL1 and TBL2. You need to minimize the amount of time required for the two queries to return records from the tables. What should you do?

- A. Create clustered indexes on TBL1 and TBL2.
- B. Create a clustered index on TBL1 Create a nonclustered index on tbl2 and add the most frequently queried columns as included columns.
- C. Create a nonclustered index on tbl2 only.
- D. Create unique constraints on both TBL1 and TBL2. Create a partitioned view that combines columns from TBL1 and TBL2.
- E. Drop existing indexes on TBL1 and then create a clustered columnstore inde
- F. Create a nonclustered columnstore index on TBL1. Create a nonclustered index on TBL2.
- G. Drop existing indexes on TBL1 and then create a cluwered columnstore inde
- H. Create a nonclustered columnstore index on TBL1. Make no changes to TBL2.
- I. Create check constraints on both TBL1 and tbl2. Create a partitioned view that combines columns from TBL1 and tbl2.
- J. Create an indexed view that combines columns from TBL1 and TBL2.

Answer: F

NEW QUESTION 70

Background
 You have a database named HR1 that includes a table named Employee.
 You have several read-only, historical reports that contain regularly changing totals. The reports use multiple queries to estimate payroll expenses. The queries run concurrently. Users report that the payroll estimate reports do not always run. You must monitor the database to identify issues that prevent the reports from running.
 You plan to deploy the application to a database server that supports other applications. You must minimize the amount of storage that the database requires.
 Employee Table
 You use the following Transact-SQL statements to create, configure, and populate the Employee table:

```

CREATE TABLE dbo.Employee
  (EmployeeId INT PRIMARY KEY,
  LastName varchar(50),
  FirstName varchar(50),
  DepartmentId int,
  HireDate datetime,
  TerminationDate datetime,
  SupervisorId int,
  CostCenterNumber int,
  EmployeeStatus int,
  EmployeePayRate int)
) GO
CREATE INDEX IX_1 on dbo.Employee (LastName, FirstName, DepartmentId) INCLUDE (HireDate)
CREATE INDEX IX_2 on dbo.Employee (LastName) INCLUDE (EmployeeId, FirstName, DepartmentId)
CREATE INDEX IX_3 on dbo.Employee (LastName, FirstName) INCLUDE (DepartmentId)
CREATE INDEX IX_4 on dbo.Employee (LastName, FirstName) INCLUDE (HireDate, DepartmentId)
GO

INSERT INTO Employee (EmployeeID, LastName, CostCenterNumber) VALUES(1001, 'Employee A', 3001001)
INSERT INTO Employee (EmployeeID, LastName, CostCenterNumber) VALUES(1002, 'Employee B', 3001001)
GO

```

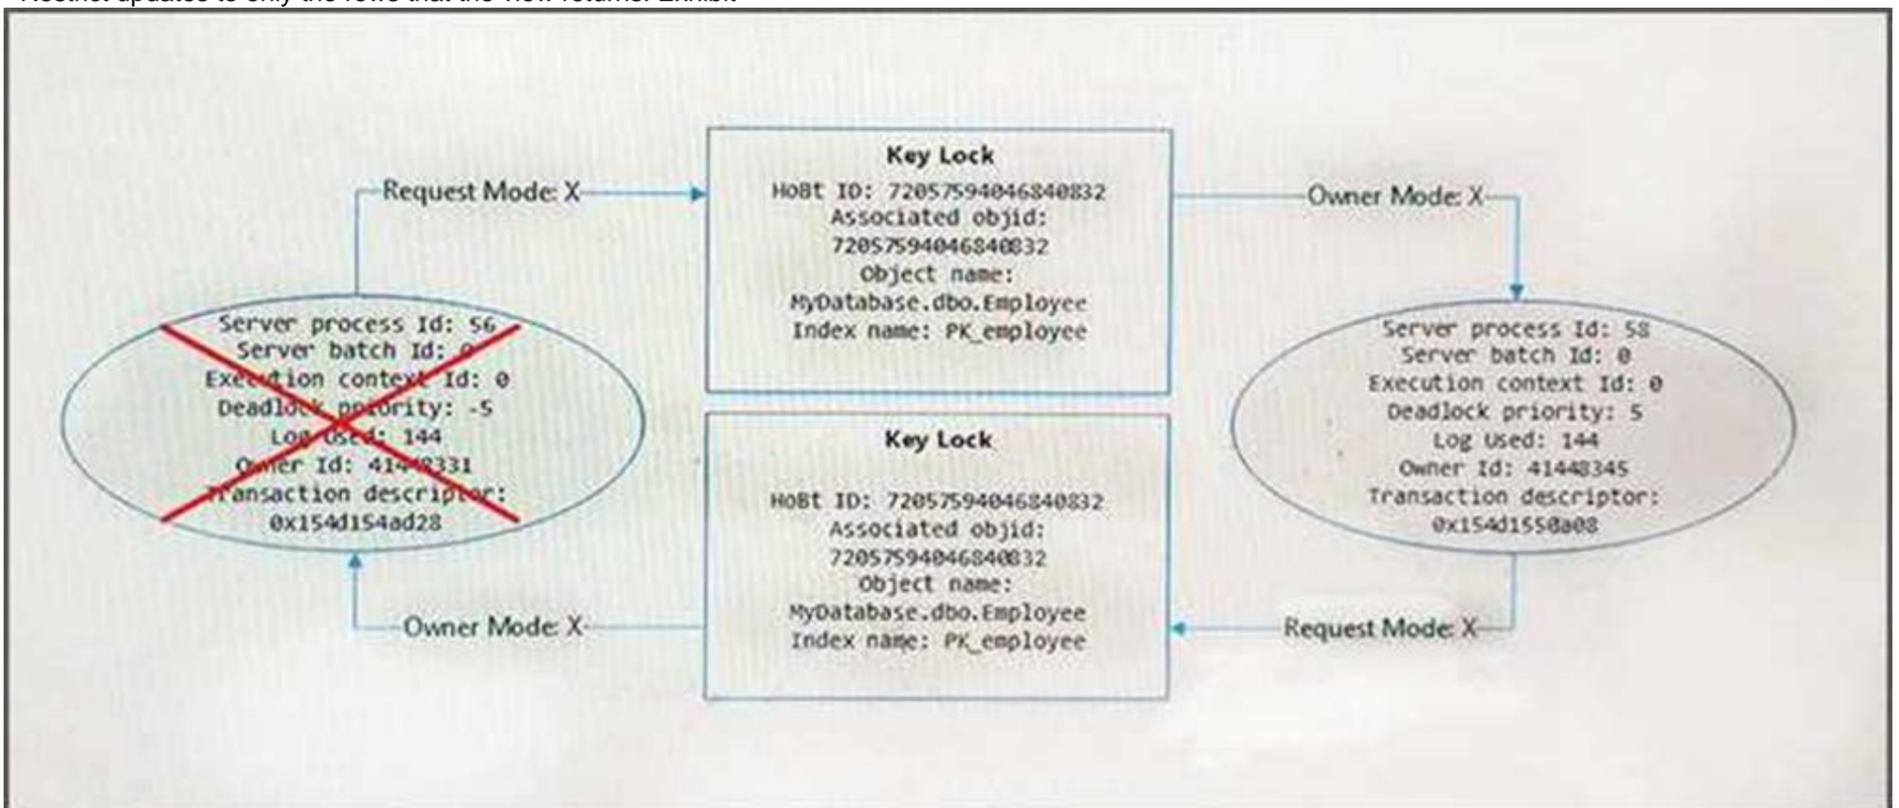
Application

You have an application that updates the Employees table. The application calls the following stored procedures simultaneously and asynchronously:

- UspA: This stored procedure updates only the EmployeeStatus column.
- UspB: This stored procedure updates only the EmployeePayRate column.

The application uses views to control access to data. Views must meet the following requirements:

- Allow user access to all columns in the tables that the view accesses.
- Restrict updates to only the rows that the view returns. Exhibit



Users must only be able to modify data in the Employee table by using the vwEmployee view. You must prevent users from viewing the view definition in catalog views.

You need to identify the view attribute to use when creating vwEmployee. In the table below, identify the attributes that you must use.

NOTE: Make only one selection in each column.

Answer Area

View creation attribute	Update restriction	Restrict access to definition
SCHEMABINDING	<input type="radio"/>	<input type="radio"/>
VIEW_METADATA	<input type="radio"/>	<input type="radio"/>
ENCRYPTION	<input type="radio"/>	<input type="radio"/>
CHECK OPTION	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation: References: <https://msdn.microsoft.com/en-us/library/ms187956.aspx>

NEW QUESTION 73

You have a reporting application that uses a table named Table1. You deploy a new batch update process to perform updates to Table1.

The environment is configured with the following properties:

The database is configured with the default isolation setting.

The application and process use the default transaction handling.

You observe the application cannot access any rows that are in use by the process.

You have the following requirements:

Ensure the application is not blocked by the process.

Ensure the application has a consistent view of the data

Ensure the application does not read dirty data.

You need to resolve the issue and meet the requirements with the least amount of administrative effort. What should you do?

- A. Enable the database for the ALLOW_SNAPSHOT_ISOLATION isolation level
- B. Modify the application for the SERIALIZABLE isolation level.
- C. Enable the database for the READ_COMMITTED_SNAPSHOT isolation level.
- D. Enable the application for the WITH (NOLOCK) hint.
- E. Enable the database for the ALLOW_SNAPSHOT_ISOLATION isolation level
- F. Modify the application and the update process for the SNAPSHOT isolation level.

Answer: B

Explanation: Snapshot isolation must be enabled by setting the ALLOW_SNAPSHOT_ISOLATION ON database option before it is used in transactions. This activates the mechanism for storing row versions in the temporary database (tempdb).

READ COMMITTED is the default isolation level for SQL Server. It prevents dirty reads by specifying that statements cannot read data values that have been modified but not yet committed by other transactions. Other transactions can still modify, insert, or delete data between executions of individual statements within the current transaction, resulting in non-repeatable reads, or "phantom" data.

NEW QUESTION 76

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

You have a database named DB1 that contains the following tables: Customer, CustomerToAccountBridge, and CustomerDetails. The three tables are part of the Sales schema. The database also contains a schema

named Website. You create the Customer table by running the following Transact-SQL statement:

```
CREATE TABLE Customer
(
    CustomerNumber int NOT NULL,
    CustomerName varchar(50) NOT NULL,
    CreateDate date NOT NULL,
    Gender bit,
    Address varchar(50)
    City varchar(50)
    State char(2),
    CustomerStatus bit NOT NULL,
    MaritalStatus bit,
    Segment varchar(5),
    CountryCode char(2),
    Birthday date,
    PostalCode char(5),
    PhoneNumber varchar(20),
    Account1 char(7),
    Account1Status bit,
    Account2 char(7),
    Account2Status bit,
    CONSTRAINT PK_Customer PRIMARY KEY CLUSTERED (CustomerNumber)
);
```

The value of the CustomerStatus column is equal to one for active customers. The value of the Account1Status and Account2Status columns are equal to one for active accounts. The following table displays selected columns and rows from the Customer table.

Customer ID	CustomerName	Gender	Account1	Account1Status	Account2	Account2Status
101	Name A	0	0001001	0	0001002	1
102	Name B	1	0002001	1	0002002	0
103	Name C	0	0003001	1	0003002	1

You plan to create a view named Website.Customer and a view named Sales.FemaleCustomers. Website.Customer must meet the following requirements:

- * Allow users access to the CustomerName and CustomerNumber columns for active customers.
- * Allow changes to the columns that the view references. Modified data must be visible through the view.
- * Prevent the view from being published as part of Microsoft SQL Server replication. Sales.FemaleCustomers must meet the following requirements:
- * Allow users access to the CustomerName, Address, City, State and PostalCode columns.
- * Prevent changes to the columns that the view references.
- * Only allow updates through the views that adhere to the view filter.

You have the following stored procedures: spDeleteCustAcctRelationship and spUpdateCustomerSummary. The spUpdateCustomerSummary stored procedure was created by running the following Transact-SQL statement:

```
CREATE PROCEDURE uspUpdateCustomerSummary
@CustomerId INT
AS
BEGIN
    SET NOCOUNT on;
    UPDATE CustomerDetails SET TotalDepositAccountCount = TotalDepositAccountCount + 1 WHERE CustomerID = @CustomerId;
    BEGIN TRAN;
    BEGIN TRY
        UPDATE CustomerDetails SET TotalAccountCount = TotalAccountCount + 1 WHERE CustomerID = @CustomerId;
    END TRY
    BEGIN CATCH
        IF @@TRANCOUNT > 0
            ROLLBACK TRAN;
    END CATCH
    IF @@TRANCOUNT > 0
        COMMIT TRAN;
```

You run the spUpdateCustomerSummary stored procedure to make changes to customer account summaries. Other stored procedures call the spDeleteCustAcctRelationship to delete records from the CustomerToAccountBridge table.

You must update the design of the Customer table to meet the following requirements.

- * You must be able to store up to 50 accounts for each customer.
- * Users must be able to retrieve customer information by supplying an account number.
- * Users must be able to retrieve an account number by supplying customer information.

Which three Transact-SQL segments should you use to develop the solution? To answer, move the appropriate code blocks from the list of code blocks to the answer area and arrange them in the correct order.

Transact-SQL segments

Answer Area

```
WHERE CustomerStatus <> 1;

ON Customer(CreateDate, CustomerStatus) INCLUDE
(Segment, ContryCode, PhoneNumber)

WHERE CustomerID = 1;

CREATE NONCLUSTERED INDEX IX_Customer_CreatedDate

ON Customer(CreateDate) INCLUDE (Segment, Country-
Code, PhoneNumber)

CREATE CLUSTERED INDEX IX_Customer_CreatedDate
```



Answer:

Explanation: Box 1: Clustered Index

With the same size of keys, the nonclustered indexes need more space than clustered indexes. Box 2, Box 3: Include the CustomerStatus column in the index, and only when CustomerStatus not equal to 1 (the active customers).

References:

<http://www.sqlserverlogexplorer.com/overview-of-cluster-and-noncluster-index/>

NEW QUESTION 79

You have a table named Person.Address that includes the following columns:

- AddressID
- AddressLine1
- AddressLine2
- City
- StateProvinceID
- PostalCode
- RowGuid
- ModifiedDate

You need to create a nonclustered index on PostalCode named IX_Address_PostalCode that uses the following included columns:

- AddressLine1
- AddressLine2
- City
- StateProvinceID

How should you complete the Transact-SQL statement? To answer, select the appropriate Transact_SQL segments in the answer are.

Answer Area

```
CREATE NONCLUSTERED [ ] IX_Address_PostalCode
[ ]
ON
INCLUDE
INDEX
WHERE
(Person.Address (PostalCode
(AddressLine1, AddressLine2, City, StateProvinceID)
[ ]
ON
INCLUDE
INDEX
WHERE
```

Answer:

Explanation: Box 1: INDEX

Box 2: ON

Box 3: INCLUDE

INCLUDE (column [,... n]) specifies the non-key columns to be added to the leaf level of the nonclustered index. The nonclustered index can be unique or non-unique.

References:

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

NEW QUESTION 80

Note: The question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other question in the series. Information and details provided in a question apply only to that question. You have a database named DB1. The database does not use a memory-optimized filegroup. The database contains a table named Table1. The table must support the following workloads:

Workload	Type	Description
Reporting	Existing	The reporting workload must scan most of the records in the table to aggregate on a number of columns. A clustered columnstore index is already created on the table to support this workload.
OLTP	New	The OLTP workload must support 3,000 transactions per second. Rows are identified by using two columns. The filter is variant on one of the two columns while constant on the other. Only a small number of records with a few columns are returned by the query.

You need to add the most efficient index to support the new OLTP workload, while not deteriorating the existing Reporting query performance. What should you do?

- A. Create a clustered index on the table.
- B. Create a nonclustered index on the table.
- C. Create a nonclustered filtered index on the table.
- D. Create a clustered columnstore index on the table.
- E. Create a nonclustered columnstore index on the table.
- F. Create a hash index on the table.

Answer: C

Explanation: A filtered index is an optimized nonclustered index, especially suited to cover queries that select from a well-defined subset of data. It uses a filter predicate to index a portion of rows in the table. A well-designed filtered index can improve query performance, reduce index maintenance costs, and reduce index storage costs compared with full-table indexes. References: [https://technet.microsoft.com/en-us/library/cc280372\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/cc280372(v=sql.105).aspx)

NEW QUESTION 81

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a database named DB1 that includes a table named Sales-Orders. You grant a user named User1 select permissions on the sales schema. You need to ensure that User1 can select data from the sales.orders table without specifying the schema name in any Transact-SQL statements.

Solution: You create a stored procedure in the sales schema that selects the data from the sales.orders table. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 83

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questions sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have a 3-TB database. The database server has 64 CPU cores. You plan to migrate the database to Microsoft Azure SQL Database.

You need to select the service tier for the Azure SQL database. The solution must meet or exceed the current processing capacity. Solution: You select the Premium service tier. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation: Premium service is required for 3 TB of storage. Single database DTU and storage limits

	Basic	Standard	Premium
Maximum storage size	2 GB	1 TB	4 TB
Maximum DTUs	5	3000	4000

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tiers-dtu>

NEW QUESTION 84

You are developing a stored procedure with the following requirements:

- *Accepts an integer as input and inserts the value into a table.
- *Ensures new transactions are committed as part of the outer transactions.
- *Preserves existing transactions if the transaction in the procedure fails.
- *If the transaction in the procedure fails, rollback the transaction.

How should you complete the procedure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```

CREATE PROCEDURE Procedure1
    @Value INT
AS BEGIN
    SET XACT_ABORT ON
    DECLARE @TranCount INT
    SET @TranCount = @@TRANCOUNT
    IF 
        XACT_STATE() = -1
        XACT_STATE() = 0
        @@TRANCOUNT = 1
    SAVE TRANSACTION SavePoint
    ELSE
        BEGIN TRANSACTION
        BEGIN TRY
            INSERT Table1 (Key1) VALUES (@Value)
            
            XACT_STATE() = -1
            XACT_STATE() = 0
            @@TRANCOUNT = 1
            COMMIT
        END TRY
        BEGIN CATCH
            IF @TranCount = 0
                ROLLBACK TRANSACTION
        ELSE IF 
            XACT_STATE() = 0
            XACT_STATE() <> -1
            @@TRANCOUNT = 1
            ROLLBACK TRANSACTION SavePoint
        END CATCH
    END

```

Answer:

Explanation:

```

SET XACT_ABORT ON
DECLARE @TranCount INT
SET @TranCount = @@TRANCOUNT
IF 
    SAVE TRANSACTION SavePoint
ELSE
    BEGIN TRANSACTION
    BEGIN TRY
        INSERT Table1 (Key1) VALUES (@Value)
        IF 
            COMMIT
    END TRY
    BEGIN CATCH
        IF @TranCount = 0
            ROLLBACK TRANSACTION
        ELSE IF 
            ROLLBACK TRANSACTION SavePoint
    END CATCH
END

```

NEW QUESTION 86

Note: this question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in the series. Information and details provided in a question apply only to that question.

You are developing an application to track customer sales.

You need to create a database object that meets the following requirements:

- Return a value of 0 if data is inserted successfully into the Customers table.
- Return a value of 1 if data is not inserted successfully into the Customers table.
- Support logic that is written by using managed code.
- Support TRY...CATCH error handling. What should you create?

- A. extended procedure
- B. CLR procedure
- C. user-defined procedure
- D. DML trigger
- E. scalar-valued function
- F. table-valued function

Answer: D

Explanation: DML triggers is a special type of stored procedure that automatically takes effect when a data manipulation language (DML) event takes place that affects the table or view defined in the trigger. DML events include INSERT, UPDATE, or DELETE statements. DML triggers can be used to enforce business rules and data integrity, query other tables, and include complex Transact-SQL statements.

References: <https://msdn.microsoft.com/en-us/library/ms178110.aspx>

NEW QUESTION 87

You are analyzing the performance of a database environment.

You suspect there are several missing indexes in the current database.

You need to return a prioritized list of the missing indexes on the current database.

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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.

Transact-SQL segments

- sys.dm_db_missing_index_group_stats
- sys.dm_db_missing_index_details
- sys.dm_db_missing_index_stats
- sql_handle
- plan_handle
- group handle

Answer Area

```

SELECT so.name
      , (avg_total_user_cost * avg_user_impact) * (user_seeks + user_scans)
as Impact
      , mid.equality_columns
      , mid.inequality_columns
      , mid.included_columns

FROM      Transact-SQL statement AS stats
INNER JOIN sys.dm_db_missing_index_groups AS mig ON stats.group_handle =
mig.index_group_handle
INNER JOIN sys.dm_db_missing_index_details AS mid ON mig.index_handle =
mid.index_handle
INNER JOIN sys.objects so WITH (nolock) ON mid.object_id = so.object_id
WHERE stats.group_handle IN (

SELECT TOP (5000) Transact-SQL statement

FROM      Transact-SQL statement WITH (nolocks)
ORDER BY (avg_total_user_cost * avg_user_impact) * (user_seeks +
user_scans) DESC)
    
```

Answer:

Explanation: Box 1: sys.db_db_missing_index_group_stats

The sys.db_db_missing_index_group_stats table include the required columns for the main query: avg_total_user_cost, avg_user_impact, user_seeks, and user_scans.

Box 2: group_handle

Example: The following query determines which missing indexes comprise a particular missing index group, and displays their column details. For the sake of this example, the missing index group handle is 24.

```

SELECT migs.group_handle, mid.*
FROM sys.dm_db_missing_index_group_stats AS migs INNER JOIN sys.dm_db_missing_index_groups AS mig ON (migs.group_handle =
mig.index_group_handle) INNER JOIN sys.dm_db_missing_index_details AS mid ON (mig.index_handle = mid.index_handle)
WHERE migs.group_handle = 24;
    
```

Box 3: sys.db_db_missing_index_group_stats

The sys.db_db_missing_index_group_stats table include the required columns for the subquery: avg_total_user_cost and avg_user_impact.

Example: Find the 10 missing indexes with the highest anticipated improvement for user queries

The following query determines which 10 missing indexes would produce the highest anticipated cumulative improvement, in descending order, for user queries.

```

SELECT TOP 10 *
FROM sys.dm_db_missing_index_group_stats
ORDER BY avg_total_user_cost * avg_user_impact * (user_seeks + user_scans)DESC;
    
```

NEW QUESTION 88

Note: this question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in the series. Information and details provided in a question apply only to that question.

You are developing an application to track customer sales.

You need to create a database object that meets the following requirements:

- Launch when table data is modified.
- Evaluate the state a table before and after a data modification and take action based on the difference.

- Prevent malicious or incorrect table data operations.
 - Prevent changes that violate referential integrity by cancelling the attempted data modification.
 - Run managed code packaged in an assembly that is created in the Microsoft.NET Framework and located into Microsoft SQL Server.
- What should you create?

- A. extended procedure
- B. CLR procedure
- C. user-defined procedure
- D. DML trigger
- E. scalar-valued function
- F. table-valued function

Answer: B

Explanation: You can create a database object inside SQL Server that is programmed in an assembly created in the Microsoft .NET Framework common language runtime (CLR). Database objects that can leverage the rich programming model provided by the CLR include DML triggers, DDL triggers, stored procedures, functions, aggregate functions, and types.

Creating a CLR trigger (DML or DDL) in SQL Server involves the following steps:

Define the trigger as a class in a .NETFramework-supported language. For more information about how to program triggers in the CLR, see CLR Triggers. Then, compile the class to build an assembly in the .NET Framework using the appropriate language compiler.

Register the assembly in SQL Server using the CREATE ASSEMBLY statement. For more information about assemblies in SQL Server, see Assemblies (Database Engine). Create the trigger that references the registered assembly.

References: <https://msdn.microsoft.com/en-us/library/ms179562.aspx>

NEW QUESTION 89

You have the following stored procedure:

```
CREATE PROCEDURE AddNextNumber @Number INT
AS
BEGIN
    SET ANSI_DEFAULTS ON
    INSERT INTO Numbers (Number) VALUES (@Number)
END
```

The Numbers table becomes unavailable when you run the stored procedure. The stored procedure obtains an exclusive lock on the table and does not release the lock.

What are two possible ways to resolve the issue? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Remove the implicit transaction and the SET ANSI_DEFAULTS ON statement.
- B. Set the ANSI_DEFAULT statement to OFF and add a COMMIT TRANSACTION statement after the INSERT statement.
- C. Add a COMMIT TRANSACTION statement after the INSERT statement.
- D. Remove the SET ANSI DEFAULTS ON statement.

Answer: B

NEW QUESTION 90

Note: This question is part of a series of questions that present the same scenario. Each question in this series contains a unique solution. Determine whether the solution meets the stated goals.

You are developing a new application that uses a stored procedure. The stored procedure inserts thousands of records as a single batch into the Employees table. Users report that the application response time has worsened since the stored procedure was updated. You examine disk-related performance counters for the Microsoft SQL Server instance and observe several high values that include a disk performance issue. You examine wait statistics and observe an unusually high WRITELOG value.

You need to improve the application response time.

Solution: You update the application to use implicit transactions when connecting to the database. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: References:

<http://sqltouch.blogspot.co.za/2013/05/writelog-waittype-implicit-vs-explicit.html>

NEW QUESTION 93

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questions sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a 3-TB database. The database server has 64 CPU cores. You plan to migrate the database to Microsoft Azure SQL Database.

You need to select the service tier for the Azure SQL database. The solution must meet or exceed the current processing capacity.

Solution: You select the Basic service tier. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

	Basic	Standard	Premium
Maximum storage size	2 GB	1 TB	4 TB
Maximum DTUs	5	3000	4000

Premium service is required for 3 TB of storage. Single database DTU and storage limits
 References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tiers-dtu>

NEW QUESTION 94

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
 You have a database that contains the following tables: BlogCategory, BlogEntry, ProductReview, Product, and SalesPerson. The tables were created using the following Transact SQL statements:

```
CREATE TABLE BlogCategory
(
    CategoryID int NOT NULL PRIMARY KEY,
    CategoryName nvarchar (20)
);

CREATE TABLE BlogEntry
(
    Entry int NOT PRIMARY KEY,
    Entrytitle nvarchar (50),
    Category int NOT NULL FOREIGN KEY REFERENCES BlogCategory
(CategoryID)
);

CREATE TABLE dbo.ProductReview
(
    ProductReviewID IDENTITY(1,1) PRIMARY KEY,
    Product int NOT NULL,
    Review varchar (1000) NOT NULL
);

CREATE TABLE dbo.Product
(
    ProductID int Identity(1,1) PRIMARY KEY,
    Name varchar(1000) NOT NULL
);

CREATE TABLE dbo.SalesPerson
(
    SalesPersonID int IDENTITY(1,1) PRIMARY KEY,
    Name varchar (1000) NOT NULL,
    SalesID Money
)
```

You must modify the ProductReview Table to meet the following requirements:

- * The table must reference the ProductID column in the Product table
- * Existing records in the ProductReview table must not be validated with the Product table.
- * Deleting records in the Product table must not be allowed if records are referenced by the ProductReview table.
- * Changes to records in the Product table must propagate to the ProductReview table.

You also have the following database tables: Order, ProductTypes, and SalesHistory, The transact-SQL statements for these tables are not available.

You must modify the Orders table to meet the following requirements:

- * Create new rows in the table without granting INSERT permissions to the table.
- * Notify the sales person who places an order whether or not the order was completed.

You must add the following constraints to the SalesHistory table:

- * a constraint on the SaleID column that allows the field to be used as a record identifier
- * a constant that uses the ProductID column to reference the Product column of the ProductTypes table
- * a constraint on the CategoryID column that allows one row with a null value in the column
- * a constraint that limits the SalePrice column to values greater than four

Finance department users must be able to retrieve data from the SalesHistory table for sales persons where the value of the SalesYTD column is above a certain threshold.

You plan to create a memory-optimized table named SalesOrder. The table must meet the following requirements:

- * The table must hold 10 million unique sales orders.
- * The table must use checkpoints to minimize I/O operations and must not use transaction logging.
- * Data loss is acceptable.

Performance for queries against the SalesOrder table that use Where clauses with exact equality operations must be optimized.

You need to enable referential integrity for the ProductReview table.

How should you complete the relevant Transact-SQL statement? To answer? select the appropriate Transact-SQL segments in the answer area. Select two alternatives.

- A. For the first selection select: WITH CHECK
- B. For the first selection select: WITH NOCHECK
- C. For the second selection select: ON DELETE NO ACTION ON UPDATE CASCADE
- D. For the second selection select: ON DELETE CASCADE ON UPDATE CASCADE

- E. For the second selection select: ON DELETE NO ACTION ON UPDATE NO ACTION
- F. For the second selection select: ON DELETE CASCADE ON UPDATE NO ACTION

Answer: BC

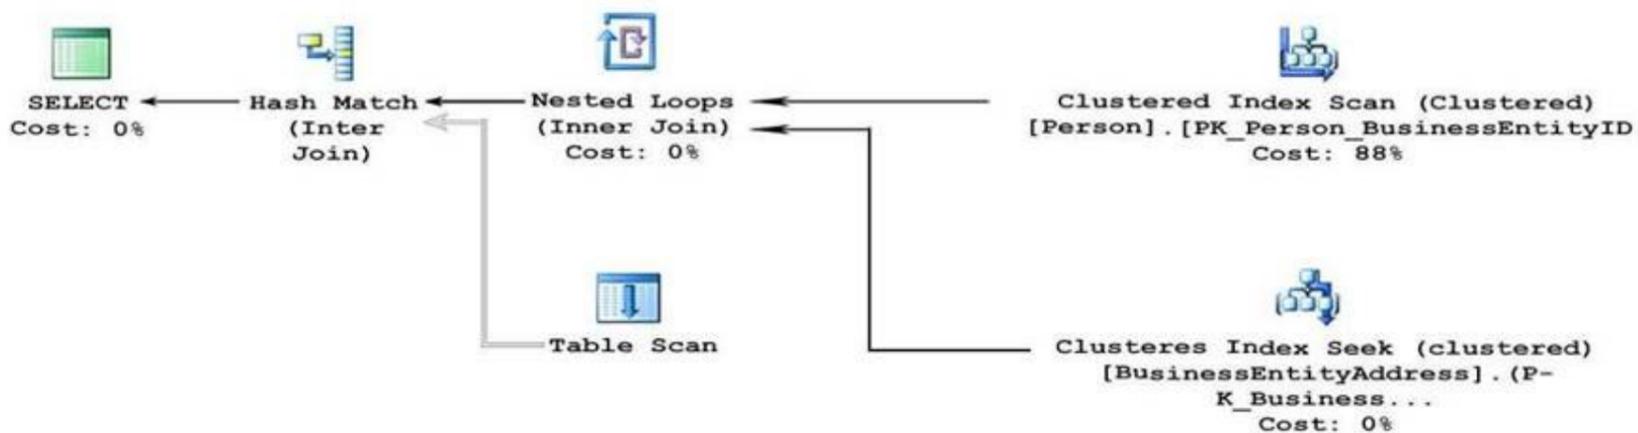
Explanation: B: We should use WITH NOCHECK as existing records in the ProductReview table must not be validated with the Product table.
 C: Deletes should not be allowed, so we use ON DELETE NO ACTION. Updates should be allowed, so we use ON DELETE NO CASCADE
 NO ACTION: the Database Engine raises an error, and the update action on the row in the parent table is rolled back.
 CASCADE: corresponding rows are updated in the referencing table when that row is updated in the parent table.
 Note: ON DELETE { NO ACTION | CASCADE | SET NULL | SET DEFAULT }
 Specifies what action happens to rows in the table that is altered, if those rows have a referential relationship and the referenced row is deleted from the parent table. The default is NO ACTION.
 ON UPDATE { NO ACTION | CASCADE | SET NULL | SET DEFAULT }
 Specifies what action happens to rows in the table altered when those rows have a referential relationship and the referenced row is updated in the parent table. The default is NO ACTION.
 Note: You must modify the ProductReview Table to meet the following requirements:
 1. The table must reference the ProductID column in the Product table
 2. Existing records in the ProductReview table must not be validated with the Product table.
 3. Deleting records in the Product table must not be allowed if records are referenced by the ProductReview table.
 4. Changes to records in the Product table must propagate to the ProductReview table. References: <https://msdn.microsoft.com/en-us/library/ms190273.aspx>
<https://msdn.microsoft.com/en-us/library/ms188066.aspx>

NEW QUESTION 99

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.
 You have a database named DB1. There is no memory-optimized filegroup in the database. You run the following query:

```
DECLARE @LastName varchar(20) = "contoso";
SELECT P.FirstName, Addr.*
From Person.Person P
JOIN Person.BusinessEntityAddress On P.BusinessEntityID =
BusinessEntityAddress.BusinessEntityID
JOIN Person.[Address] Addr ON BusinessEntityAddress.AddressID =
Addr.AddressID
WHERE P.LastName = @LastName;
```

The following image displays the execution plan the query optimizer generates for this query:



Users frequently run the same query with different values for the local variable @lastName. The table named Person is persisted on disk. You need to create an index on the Person.Person table that meets the following requirements:
 - All users must be able to benefit from the index.
 - FirstName must be added to the index as an included column. What should you do?

- A. Create a clustered index on the table.
- B. Create a nonclustered index on the table.
- C. Create a nonclustered filtered index on the table.
- D. Create a clustered columnstore index on the table.
- E. Create a nonclustered columnstoreindex on the table.
- F. Create a hash index on the table.

Answer: B

Explanation: By including nonkey columns, you can create nonclustered indexes that cover more queries. This is because the nonkeycolumns have the following benefits:
 They can be data types not allowed as index key columns.
 They are not considered by the Database Engine when calculating the number of index key columns or index key size.

NEW QUESTION 102

Note: This question is part of a series of questions that use the same or similar answer choices. An Answer choice may be correct for more than one question in the series. Each question independent of the other questions in this series. Information and details provided in a question apply only to that question.
 You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run during off-peak hours.
 You must monitor and optimize the SQL Server to maximize throughput, response time, and overall SQL performance.
 You need to identify previous situations where a modification has prevented queries from selecting data in tables.
 What should you do?

- A. Create a sys.dm_os_waiting_tasks query.
- B. Create a sys.dm_exec_sessions query.
- C. Create a Performance Monitor Data Collector Set.
- D. Create a sys.dm_os_memory_objects query.
- E. Create a sp_configure 'max server memory' query.
- F. Create a SQL Profiler trace.
- G. Create a sys.dm_os_wait_stats query.
- H. Create an Extended Event.

Answer: G

Explanation: sys.dm_os_wait_stats returns information about all the waits encountered by threads that executed. You can use this aggregated view to diagnose performance issues with SQL Server and also with specific queries and batches.

NEW QUESTION 106

You have a database that contains both disk-based and memory-optimized tables.
 You need to create two modules. The modules must meet the requirements described in the following table.

Module	Requirements
Module 1	<ul style="list-style-type: none"> - must be encrypted by using the ENCRYPTON option - must support updates on both disk-based and memory-optimized tables - must support OUTPUT parameters
Module 2	<ul style="list-style-type: none"> - must access only memory-optimized tables - must support updates on memory-optimized tables - must support heavy aggregations with highest performance - must support OUTPUT parameters

Which programming object should you use for each module? To answer, select the appropriate object types in the answer area.

Answer Area

Module	Object type
Module 1	<div style="border: 1px solid gray; padding: 5px;"> <ul style="list-style-type: none"> interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function </div>
Module 2	<div style="border: 1px solid gray; padding: 5px;"> <ul style="list-style-type: none"> interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function </div>

Answer:

Explanation: Returning Data by Using OUTPUT Parameters

If you specify the OUTPUT keyword for a parameter in the procedure definition, the stored procedure can return the current value of the parameter to the calling program when the stored procedure exits.
 SQL Server stored procedures, views and functions are able to use the WITH ENCRYPTION option to disguise the contents of a particular procedure or function from discovery.

Native Compilation of Tables and Stored Procedures

In-Memory OLTP introduces the concept of native compilation. SQL Server can natively compile stored procedures that access memory-optimized tables. SQL Server is also able to natively compile memory-optimized tables. Native compilation allows faster data access and more efficient query execution than interpreted (traditional) Transact-SQL. Native

compilation of tables and stored procedures produce DLLs.

References: [https://technet.microsoft.com/en-us/library/ms187004\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms187004(v=sql.105).aspx) <https://msdn.microsoft.com/en-us/library/dn249342.aspx>

NEW QUESTION 110

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a database that contains a table named Employees. The table stores information about the employees of your company.

You need to implement the following auditing rules for the Employees table:

- Record any changes that are made to the data in the Employees table.
- Customize the data recorded by the audit operations.

Solution: You implement a check constraint on the Employees table. Does the solution meet the goal?

- A. Yes
- B. No

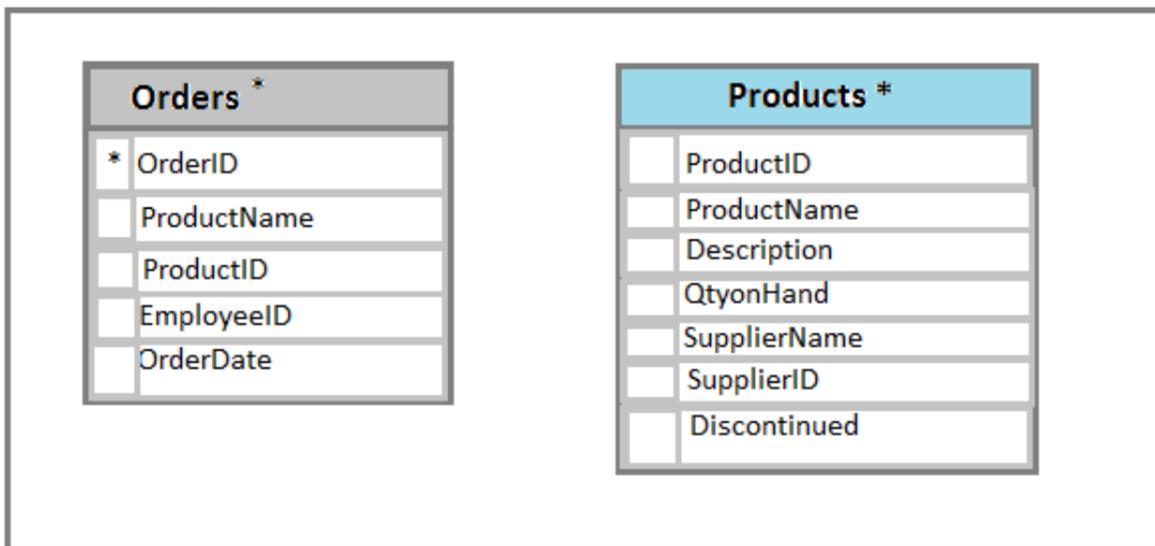
Answer: B

Explanation: Check constraints cannot be used to track changes in a table. References: <https://msdn.microsoft.com/en-us/library/bb933994.aspx>

NEW QUESTION 115

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

You have a database named Sales that contains the following database tables: Customer, Order, and Products. The Products table and the Order table are shown in the following diagram.



The customer table includes a column that stores the data for the last order that the customer placed.

You plan to create a table named Leads. The Leads table is expected to contain approximately 20,000 records. Storage requirements for the Leads table must be minimized.

You need to create triggers that meet the following requirements:

- * Optimize the performance and data integrity of the tables.
- * Provide a custom error if a user attempts to create an order for a customer that does not exist.
- * In the Customers table, update the value for the last order placed.
- * Complete all actions as part of the original transaction.

In the table below, identify the trigger types that meet the requirements.

NOTE: Make only selection in each column. Each correct selection is worth one point.

Answer Area		
Trigger type	Provide custom	Update Customer table
AFTER INSERT trigger	<input type="radio"/>	<input type="radio"/>
INSTEAD OF INSERT trigger	<input type="radio"/>	<input type="radio"/>
AFTER UPDATE trigger	<input type="radio"/>	<input type="radio"/>
INSTEAD OF UPDATE trigger	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation: INSTEAD OF INSERT triggers can be defined on a view or table to replace the standard action of the INSERT statement. AFTER specifies that the DML trigger is fired only when all operations specified in the triggering SQL statement have executed successfully.

References: [https://technet.microsoft.com/en-us/library/ms175089\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms175089(v=sql.105).aspx)

NEW QUESTION 116

You are analyzing the memory usage of a Microsoft SQL Server instance. You need to obtain the information described on the following table.

Requirement	Details
Requirement 1	Total amount of memory currently used by SQL Server
Requirement 2	Total amount of memory required by SQL Server for running processers efficiently
Requirement 3	Total amount of memory used by a process

Which performance counter should you use for each requirement? To answer, drag the appropriate performance counters to the correct requirements. Each performance counter may be used once, more than once or not at all. You may need to drag the split bat between panes or scroll to view content. NOTE: Each correct selection is worth one point.

Transact-SQL segments

- Memory: Available Bytes
- SQL Server: Memory Manager: SQL Cache Memory (KB)
- SQL Server: Buffer Manager: Page reads/sec
- SQL Server: Memory Manager: Total Server Memory (KB)
- SQL Server: Memory Manager: Target Server Memory (KB)
- SQL Server: Memory Manager: Granted Workspace Memory (KB)
- SQL Server: Memory Manager: Maximum Workspace Memory (KB)
- Process: working Set

Answer Area

Requirement	Performancr counter
Requirement 1	Performancr counter
Requirement 2	Performancr counter
Requirement 3	Performancr counter

Answer:

Explanation: Requirement1: SQL Server: Memory Manager: Total Server Memory (KB)

This counter specifies the amount of memory the server has committed using the memory manager. Requirement2: SQL Server: Memory Manager: Granted Workspace Memory (KB)

Specifies the total amount of memory currently granted to executing processes, such as hash, sort, bulk copy, and index creation operations.

Requirement3: Process: working Set

Each time a process is created, it reserves the minimum working set size for the process. The virtual memory manager attempts to keep enough memory for the minimum working set resident when the process is active, but keeps no more than the maximum size.

References:

<https://msdn.microsoft.com/en-us/library/ms190924.aspx> <https://blogs.technet.microsoft.com/askperf/2007/05/18/sql-and-the-working-set/>

NEW QUESTION 117

You maintain a Microsoft Azure SQL Database instance.

You grant User1 the SELECT and EXECUTE permissions for all objects in the dbo schema. You must create a stored procedure that allows User1 to view the following information:

details for each connection to the database

a list of all active user connections and internal tasks

You need to create the stored procedure for User1 and ensure that User1 can run the stored procedure without any error.

How should you complete the Transact-SQL statements? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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.

NOTE: More than one combination of answer choices is correct. You will receive credit for any of the correct combinations you select.

Transact-SQL segments

GRANT SELECT ON SCHEMA: :sys TO User1

GRANT VIEW DATABASE STATE TO User1

GRANT VIEW DEFINITION TO User1

SELECT = FROM sys.dm_exec_connections

SELECT = FROM sys.dm_exec_requests

SELECT = FROM sys.dm_exec_sessions

Answer Area

```

CREATE PROCEDURE ViewConnections
AS
    Transact-SQL segment
    Transact-SQL segment
GO
    Transact-SQL segment
    
```

Answer:

Explanation: Box 1: Sys.dm_exec_connections

Sys.dm_exec_connections returns information about the connections established to this instance of SQL Server and the details of each connection. Returns server wide connection information for SQL Server. Returns current database connection information for SQL Database.

Box 2: sys.dm_exec_sessions

sys.dm_exec_sessions returns one row per authenticated session on SQL Server. sys.dm_exec_sessions is a server-scope view that shows information about all active user connections and internal tasks.

Box 3: GRANT VIEW DATABASE STATE To User1

SQL Database: Requires VIEW DATABASE STATE to see all connections to the current database. VIEW DATABASE STATE cannot be granted in the master database.

NEW QUESTION 119

You are developing a database reporting solution for a table that contains 900 million rows and is 103 GB. The table is updated thousands of times a day, but data is not deleted.

The SELECT statements vary in the number of columns used and the amount of rows retrieved.

You need to reduce the amount of time it takes to retrieve data from the table. The must prevent data duplication.

Which indexing strategy should you use?

- A. a nonclustered index for each column in the table
- B. a clustered columnstore index for the table
- C. a hash index for the table
- D. a clustered index for the table and nonclustered indexes for nonkey columns

Answer: B

Explanation: Columnstore indexes are the standard for storing and querying large data warehousing fact tables. It uses column-based data storage and query processing to achieve up to 10x query performance gains in your data warehouse over traditional row-oriented storage.

A clustered columnstore index is the physical storage for the entire table.

Generally, you should define the clustered index key with as few columns as possible.

A nonclustered index contains the index key values and row locators that point to the storage location of the table data. You can create multiple nonclustered indexes on a table or indexed view. Generally, nonclustered indexes should be designed to improve the performance of frequently used queries that are not covered by the clustered index.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview?view=sql-serv>

NEW QUESTION 124

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a table that has a clustered index and a nonclustered index. The indexes use different columns from the table. You have a query named Query1 that uses the nonclustered index.

Users report that Query1 takes a long time to report results. You run Query1 and review the following statistics for an index seek operation:

Index Seek (NonClustered)

Scan a particular range of rows from a nonclustered index.

Physical Operation	Index Seek
Logical Operation	Index Seek
Actual Execution Mode	Row
Actual Number of Rows	3571454
Actual Number of Batches	0
Estimated I/O Cost	0.0093577
Estimated Operator Cost	0.0107304 (0%)
Estimated CPU Cost	0.0013727
Estimated Subtree Cost	0.0107304
Estimated Number of Executions	1
Number of Executions	8
Estimated Number of Rows	0
Estimated Row Size	19 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	100

You need to resolve the performance issue.

Solution: You update statistics for the nonclustered index. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation: We see Actual Number of Row is 3571454, while Estimated Number of Rows is 0. This indicates that the statistics are old, and need to be updated.

NEW QUESTION 129

You need to build a function that meets the following requirements:

*Returns multiple rows

*Optimizes the performance of the query within the function

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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.

Transact-SQL segments

```

RETURNS TABLE AS

BEGIN
INSERT INTO @Prices
SELECT ListPrice FROM Product
WHERE ProductId = @ProductId
RETURN
END

RETURNS @Prices TABLE (ListPrice money) AS

RETURNS [Money] AS

BEGIN
SELECT @ListPrice = ListPrice FROM Product
WHERE ProductId = @ProductId
RETURN @ListPrice
END
    
```

Answer area

```

CREATE FUNCTION dbo.GetPrice (@Produ
    
```

Transact-SQL segment

Transact-SQL segment

Answer:

Explanation:

```

BEGIN
INSERT INTO @Prices
SELECT ListPrice FROM Product
WHERE ProductId = @ProductId
RETURN
END

RETURNS @Prices TABLE (ListPrice money) AS
    
```

Answer area

```

BEGIN
INSERT INTO @Prices
SELECT ListPrice FROM Product
WHERE ProductId = @ProductId
RETURN
END

RETURN (
SELECT ListPrice FROM Product
WHERE ProductId = @ProductId
)
    
```

NEW QUESTION 132

Note: This question is part of a series of questions that present the same scenario. Each question in this series contains a unique solution. Determine whether the solution meets the stated goals.

You are developing a new application that uses a stored procedure. The stored procedure inserts thousands of records as a single batch into the Employees table. Users report that the application response time has worsened since the stored procedure was updated. You examine disk-related performance counters for the Microsoft SQL Server instance and observe several high values that include a disk performance issue. You examine wait statistics and observe an unusually high WRITELOG value.

You need to improve the application response time.

Solution: You replace the stored procedure with a user-defined function. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: References: <https://msdn.microsoft.com/en-us/library/ms345075.aspx>

NEW QUESTION 133

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-762 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-762 Product From:

<https://www.2passeasy.com/dumps/70-762/>

Money Back Guarantee

70-762 Practice Exam Features:

- * 70-762 Questions and Answers Updated Frequently
- * 70-762 Practice Questions Verified by Expert Senior Certified Staff
- * 70-762 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * 70-762 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year