

## 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**

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
```

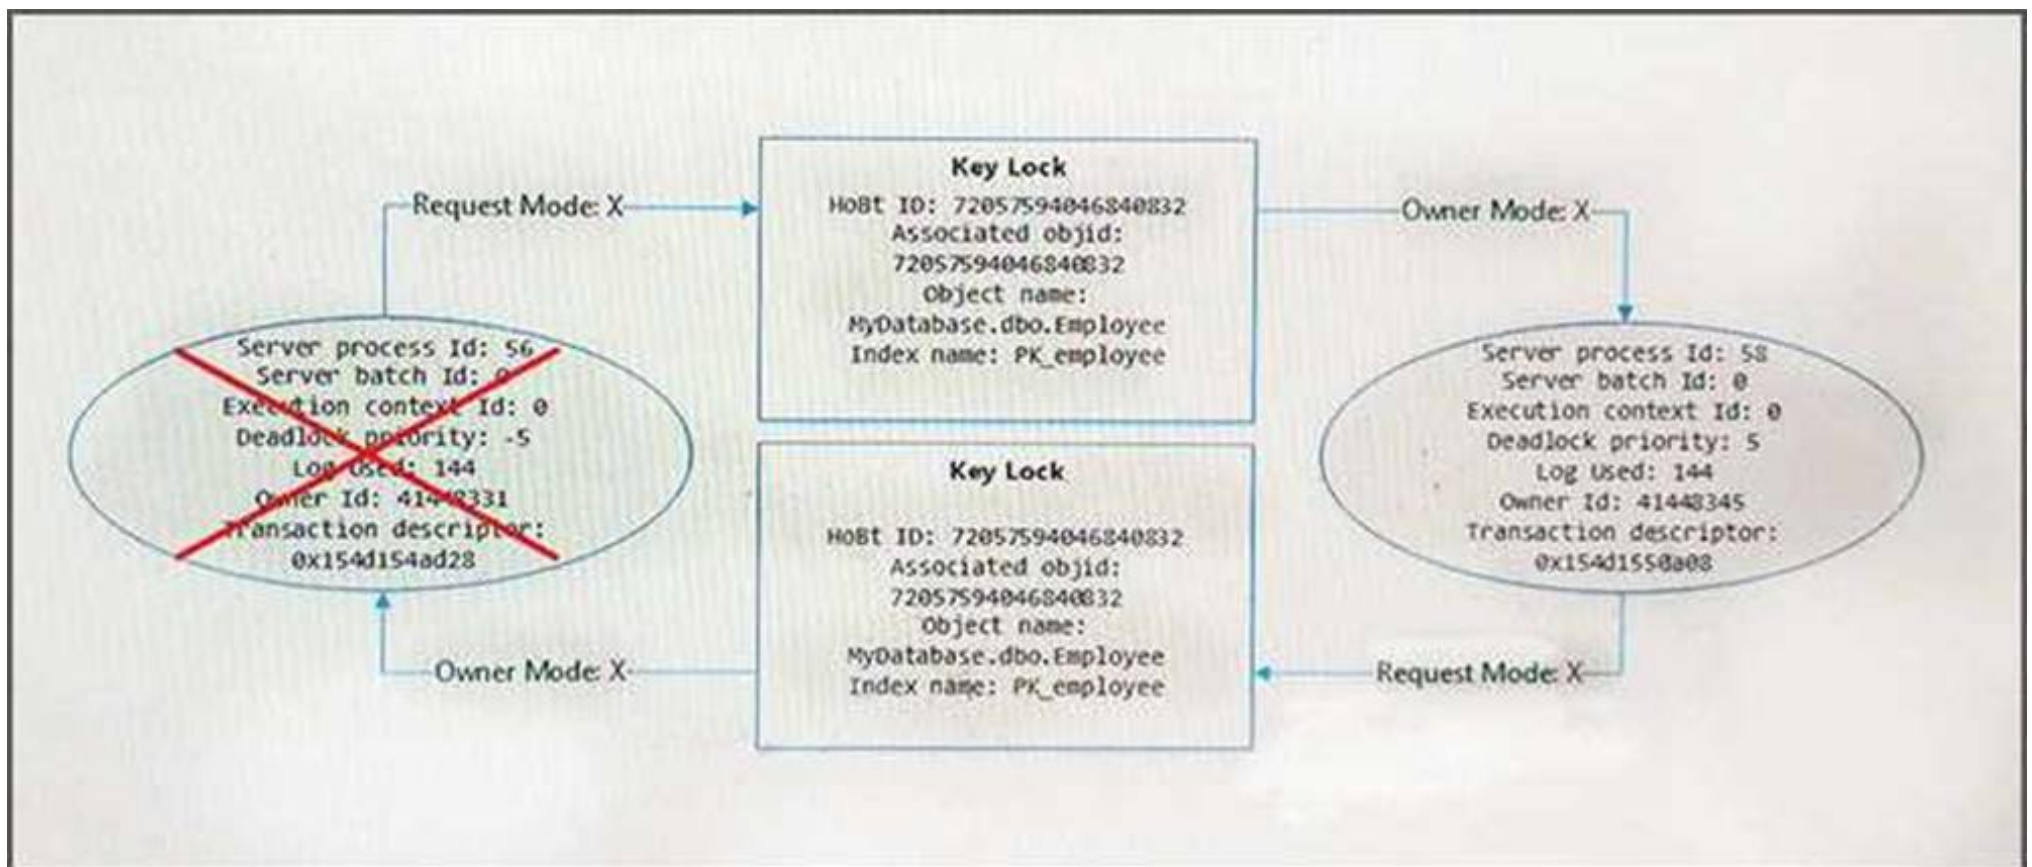
Application

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 are analyzing the performance of the database environment. You discover that locks that are held for a long period of time as the reports are generated. You need to generate the reports more quickly. The database must not use additional resources. What should you do?

- A. Update the transaction level of the report query session to READPAST.
- B. Modify the report queries to use the UNION statement to combine the results of two or more queries.
- C. Set the READ\_COMMITTED\_SNAPSHOT database option to ON.
- D. Update the transaction level of the report query session to READ UNCOMMITTED.

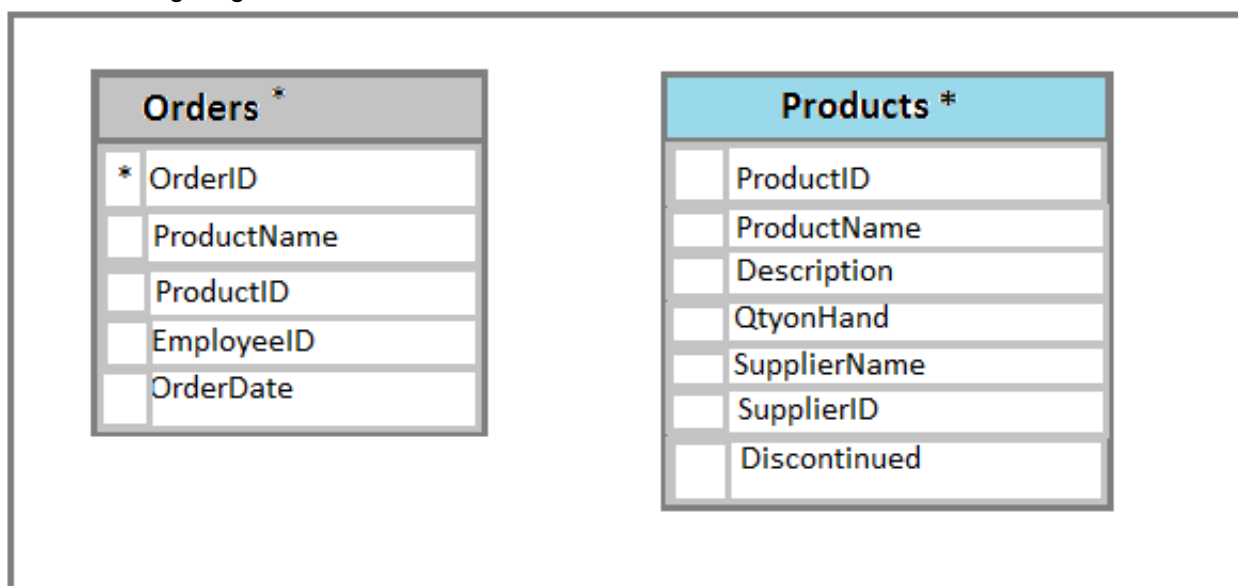
**Answer: C**

**Explanation:** References: [https://technet.microsoft.com/en-us/library/ms173763\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms173763(v=sql.105).aspx)

### NEW QUESTION 3

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 implement a stored procedure that deletes a discontinued product from the Products table. You identify the following requirements:

- \* If an open order includes a discontinued product, the records for the product must not be deleted.
  - \* The stored procedure must return a custom error message if a product record cannot be deleted. The message must identify the OrderID for the open order.
- What should you do? To answer, select the appropriate Transact-SQL segments in the answer area.



## Answer Area

### Requirement

Handle errors

### Transact-SQL segment

Try/Parse
Select @@error
Begin Tran/Rollback Tran
Try/Catch*

Display error message

ERROR MESSAGE()
PRINT
RAISERROR
RETURN

Answer:

**Explanation:** Using TRY...CATCH in Transact-SQL

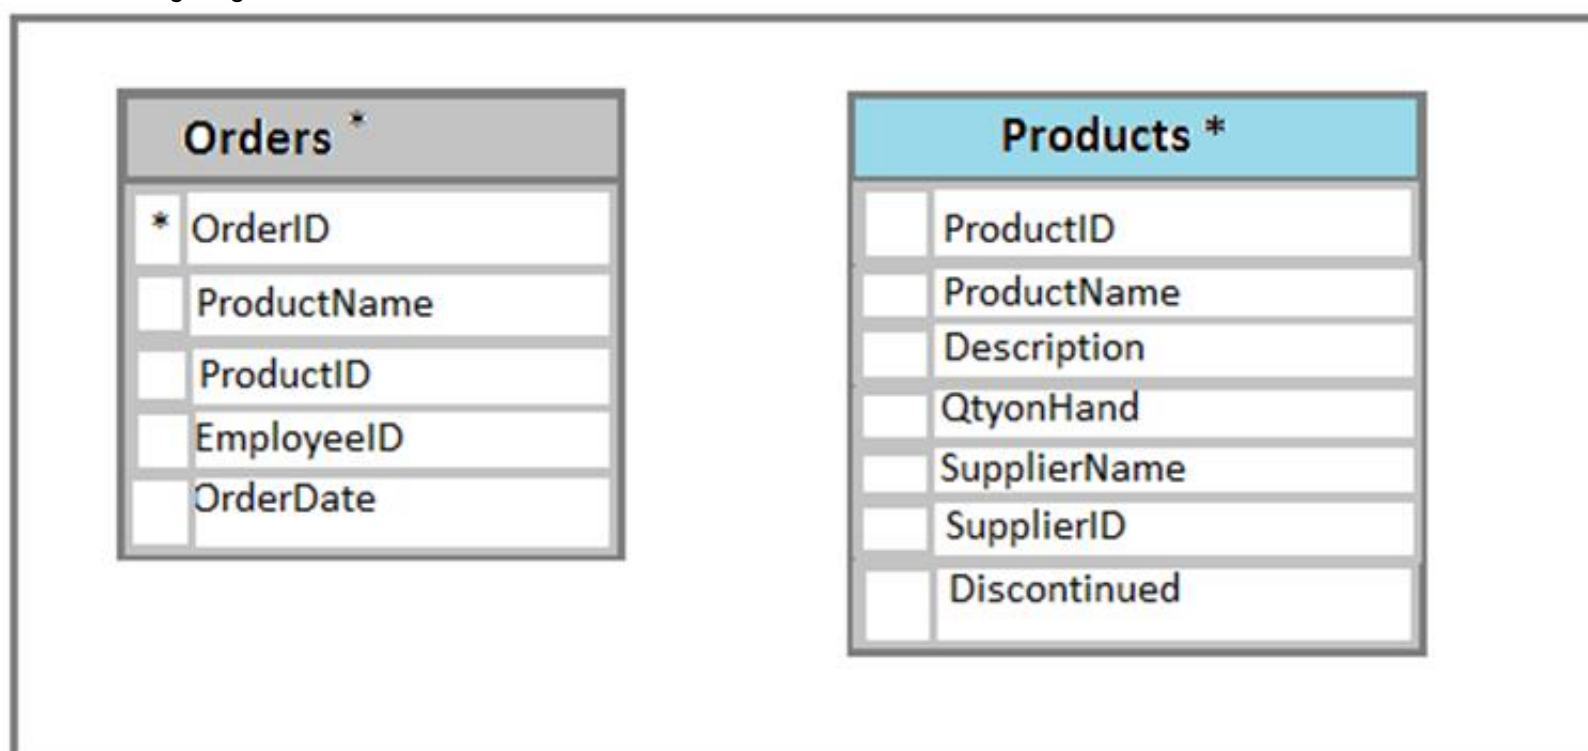
Errors in Transact-SQL code can be processed by using a TRY...CATCH construct. TRY...CATCH can use the following error function to capture error information: ERROR\_MESSAGE() returns the complete text of the error message. The text includes the values supplied for any substitutable parameters such as lengths, object names, or times.

References: [https://technet.microsoft.com/en-us/library/ms179296\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms179296(v=sql.105).aspx)

### NEW QUESTION 4

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 5

You have a database named Database1.

Users report that they experience deadlock issues- You run the sp\_readerlog stored procedure. You view the output from the Process List section as shown in the Process List exhibit. (Click the Exhibit button.)

\*\*\*Exhibit is Missing\*\*\*

You view the contents of the Resource List section as shown in the Resource List exhibit, (Click the Exhibit button.)

\*\*\*Exhibit is Missing\*\*\*

You view deadlock information as shown in the Deadlock List exhibit. (Click the Exhibit button.)

\*\*\*Exhibit is Missing\*\*\*

	Yes	No
User1's transaction has an exclusive lock on Table01.	<input type="radio"/>	<input type="radio"/>
User2's transaction uses a custom execution level.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

	Yes	No
User1's transaction has an exclusive lock on Table01.	<input checked="" type="radio"/>	<input type="radio"/>
User2's transaction uses a custom execution level.	<input type="radio"/>	<input checked="" type="radio"/>

### NEW QUESTION 6

You are designing a stored procedure for a database named obi.

The following requirements must be met during the entire execution of the stored procedure:

\*The stored procedure must only read changes that are persisted to the database.

\*select statements within the stored procedure should only show changes to the data that are made by the stored procedure

You need to configure the transaction isolation level for the stored procedure. Which Transact-SQL statement or statements should you run?

A)

```
SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED
ALTER DATABASE DB1 SET READ_COMMITTED_SNAPSHOT ON
```

B)

```
SET TRANSACTION ISOLATION LEVEL READ COMMITTED
ALTER DATABASE DB1 SET READ_COMMITTED_SNAPSHOT OFF
```

C)

```
SET TRANSACTION ISOLATION LEVEL SERIALIZABLE
```

D)

```
SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED
ALTER DATABASE SET READ_COMMITTED_SNAPSHOT OFF
```

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

Answer: A

NEW QUESTION 7

You manage a database with tables named Invoice and InvoiceDetails. Each invoice may have multiple records. Users update the InvoiceDetails table by using a .NET web application. The application retrieves records from both tables and updates the tables by running an inline update statement. Users experience slow performance when updating records in the application. The solution must meet the following requirements:

- Must use a stored procedure.
- Must not use inline update statements
- Must use a table-valued parameter.
- Must call the stored procedure to update all records. You need to optimize performance.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

- Create a stored procedure and use invDetails as a read-only input parameter.
- Create a stored procedure and use invDetails as the input parameter.
- Send invoice detail records to the stored procedure as a string value for the input parameter.
- Create a user-defined table type named invDetails like schema of InvoiceDetails table.
- Create an alias type from VARCHAR (MAX) Named invDetails.
- Send invoice detail records to the stored procedure as a ADO.NET table for the input parameter

Answer Area



Answer:

**Explanation:** Box 1: Create a user-defined table type...  
Table-valued parameters are declared by using user-defined table types. You can use table-valued parameters to send multiple rows of data to a Transact-SQL statement or a routine, such as a stored procedure or function, without creating a temporary table or many parameters.  
Box 2: ...read-only input parameter.  
Table-valued parameters must be passed as input READONLY parameters to Transact-SQL routines. Box 3:  
Example  
The following example uses Transact-SQL and shows you how to create a table-valued parameter type, declare a variable to reference it, fill the parameter list, and then pass the values to a stored procedure.  
USE AdventureWorks2012;



```

/* Create a table type. */
CREATE TYPE LocationTableType AS TABLE ( LocationName VARCHAR(50)
, CostRate INT ); GO
/* Create a procedure to receive data for the table-valued parameter. */ CREATE PROCEDURE dbo. usp_InsertProductionLocation
@TVP LocationTableType READONLY Etc.."
/* Declare a variable that references the type. */ DECLARE @LocationTVP AS LocationTableType;
/* Add data to the table variable. */
INSERT INTO @LocationTVP (LocationName, CostRate) SELECT Name, 0.00
FROM AdventureWorks2012.Person.StateProvince;
/* Pass the table variable data to a stored procedure. */ EXEC usp_InsertProductionLocation @LocationTVP; GO
References:
https://docs.microsoft.com/en-us/sql/relational-databases/tables/use-table-valued-parameters-database-engine?vi

```

### NEW QUESTION 8

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

	Isolation Level	
<div style="font-size: 4em; line-height: 1;">↑</div> <div style="writing-mode: vertical-rl; transform: rotate(180deg);">Concurrency</div>	Isolation level	<div style="font-size: 4em; line-height: 1;">↓</div> <div style="writing-mode: vertical-rl; transform: rotate(180deg);">Consistency</div>
	Isolation level	
	Isolation level	
	Isolation level	

**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 9

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 an object that meets the following requirements:

- Run managed code packaged in an assembly that was created in the Microsoft.NET Framework and uploaded in Microsoft SQL Server.
- Run within a transaction and roll back if a failure occurs.
- Run when a table is created or modified. 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
- G. DDL trigger

**Answer:** B

**Explanation:** The common language runtime (CLR) is the heart of the Microsoft .NET Framework and provides the execution environment for all .NET Framework code. Code that runs within the CLR is referred to as managed code.

With the CLR hosted in Microsoft SQL Server (called CLR integration), you can author stored procedures, triggers, user-defined functions, user-defined types, and user-defined aggregates in managed code. Because managed code compiles to native code prior to execution, you can achieve significant performance increases in some scenarios.

#### NEW QUESTION 10

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 theAccountNumber 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 10

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
```

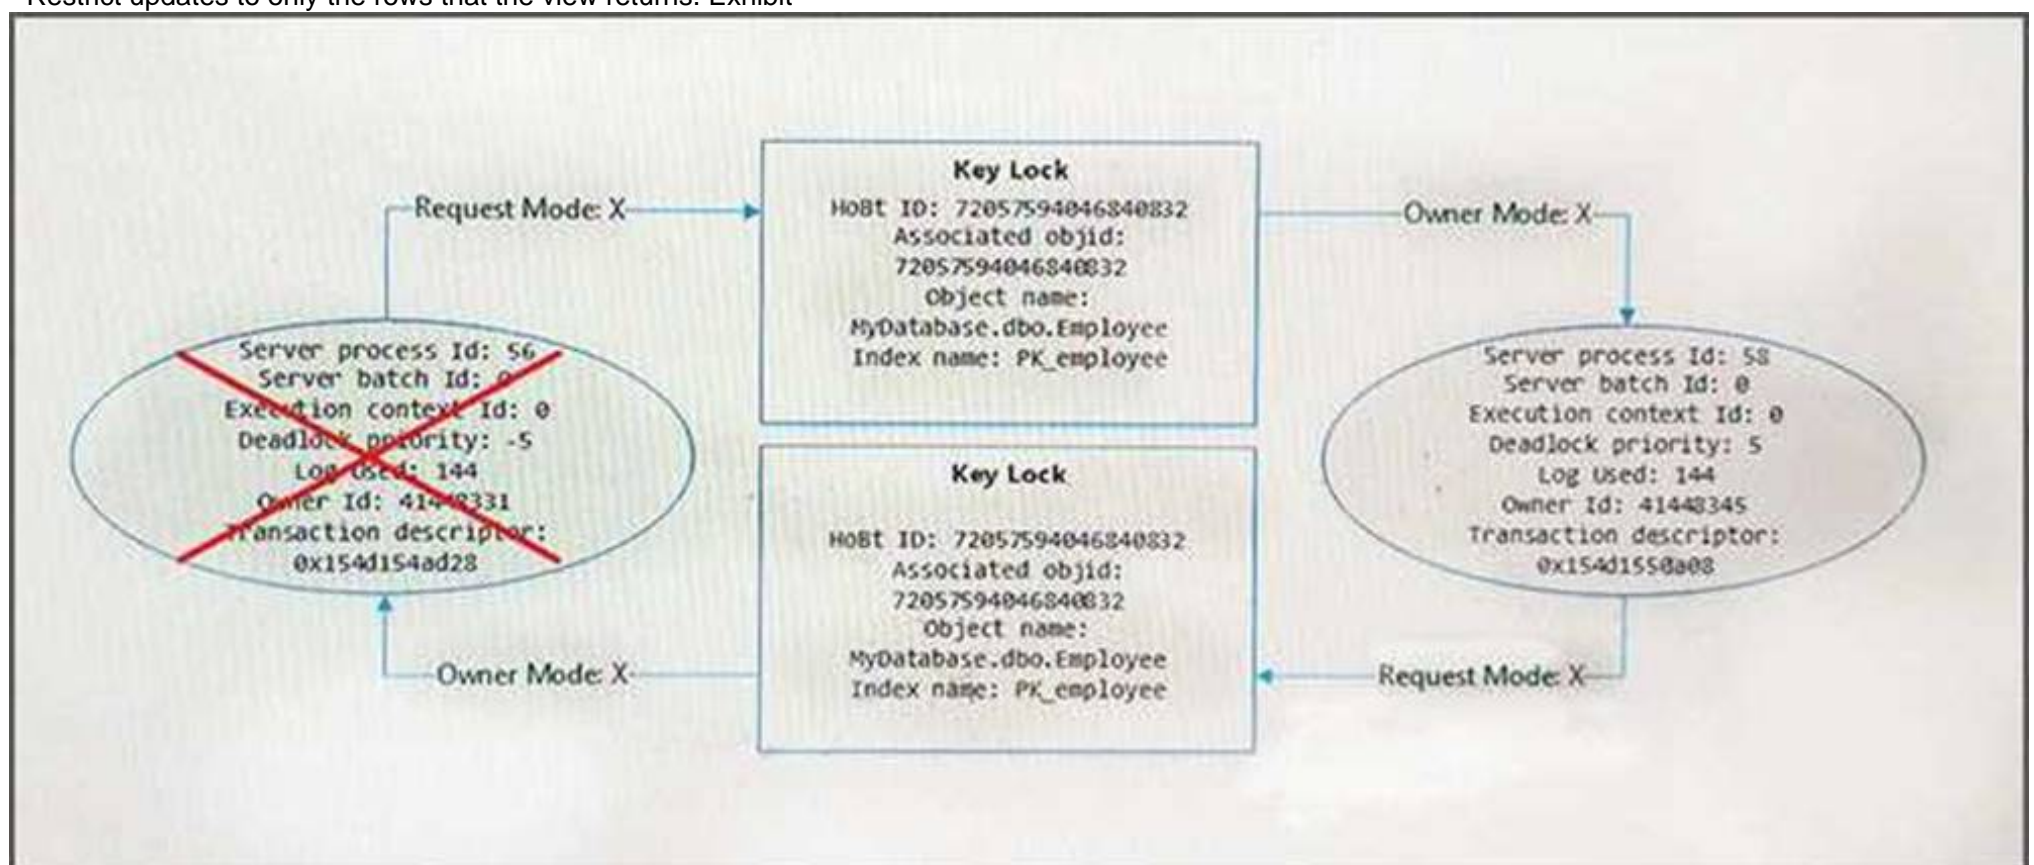
#### Application

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 if 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 15

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 17

Database users report that SELECT statements take a long time to return results. You run the following Transact-SQL statement:

```
SELECT OBJECT_NAME([object_id]) AS [object_name], d.equality_columns, d.inequality_columns, d.included_columns  
FROM sys.dm_db_missing_index_details;
```

You need to create one nonclustered index that contains all of the columns in the above table. You must minimize index

- A. CREATE NONCLUSTERED INDEX IX\_User ON Users (CountryCode, UserStatus, UserName);
- B. CREATE NONCLUSTERED INDEX IX\_User ON Users (CountryCode, UserStatus) INCLUDE (UserName);
- C. CREATE NONCLUSTERED INDEX IX\_User ON Users (CountryCode, UserName);
- D. CREATE NONCLUSTERED INDEX IX\_User ON Users (UserStatus, CountryCode) INCLUDE (UserName);

**Answer: D**

#### NEW QUESTION 20

Your company runs end-of-the-month accounting reports. While the reports run, other financial records are updated in the database.

Users report that the reports take longer than expected to run.

You need to reduce the amount of time it takes for the reports to run. The reports must show committed data only.

What should you do?

- A. Use the NOLOCK option.
- B. Execute the BDCC 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 XACT ABORT to ON.
- G. Execute the alter table t1 set (Lock\_ESCALATION = AUTO); statement.
- H. Use the output parameters.

**Answer: G**

#### NEW QUESTION 23

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. DDL trigger
- F. scalar-valued function
- G. table-valued function

**Answer: B**

**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.

A CLR trigger is a type of DDL trigger. A CLR Trigger can be either an AFTER or INSTEAD OF trigger. A CLR trigger can also be a DDL trigger. Instead of executing a Transact-SQL stored procedure, a CLR trigger executes one or more methods written in managed code that are members of an assembly created in the .NET Framework and uploaded in SQL Server.

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

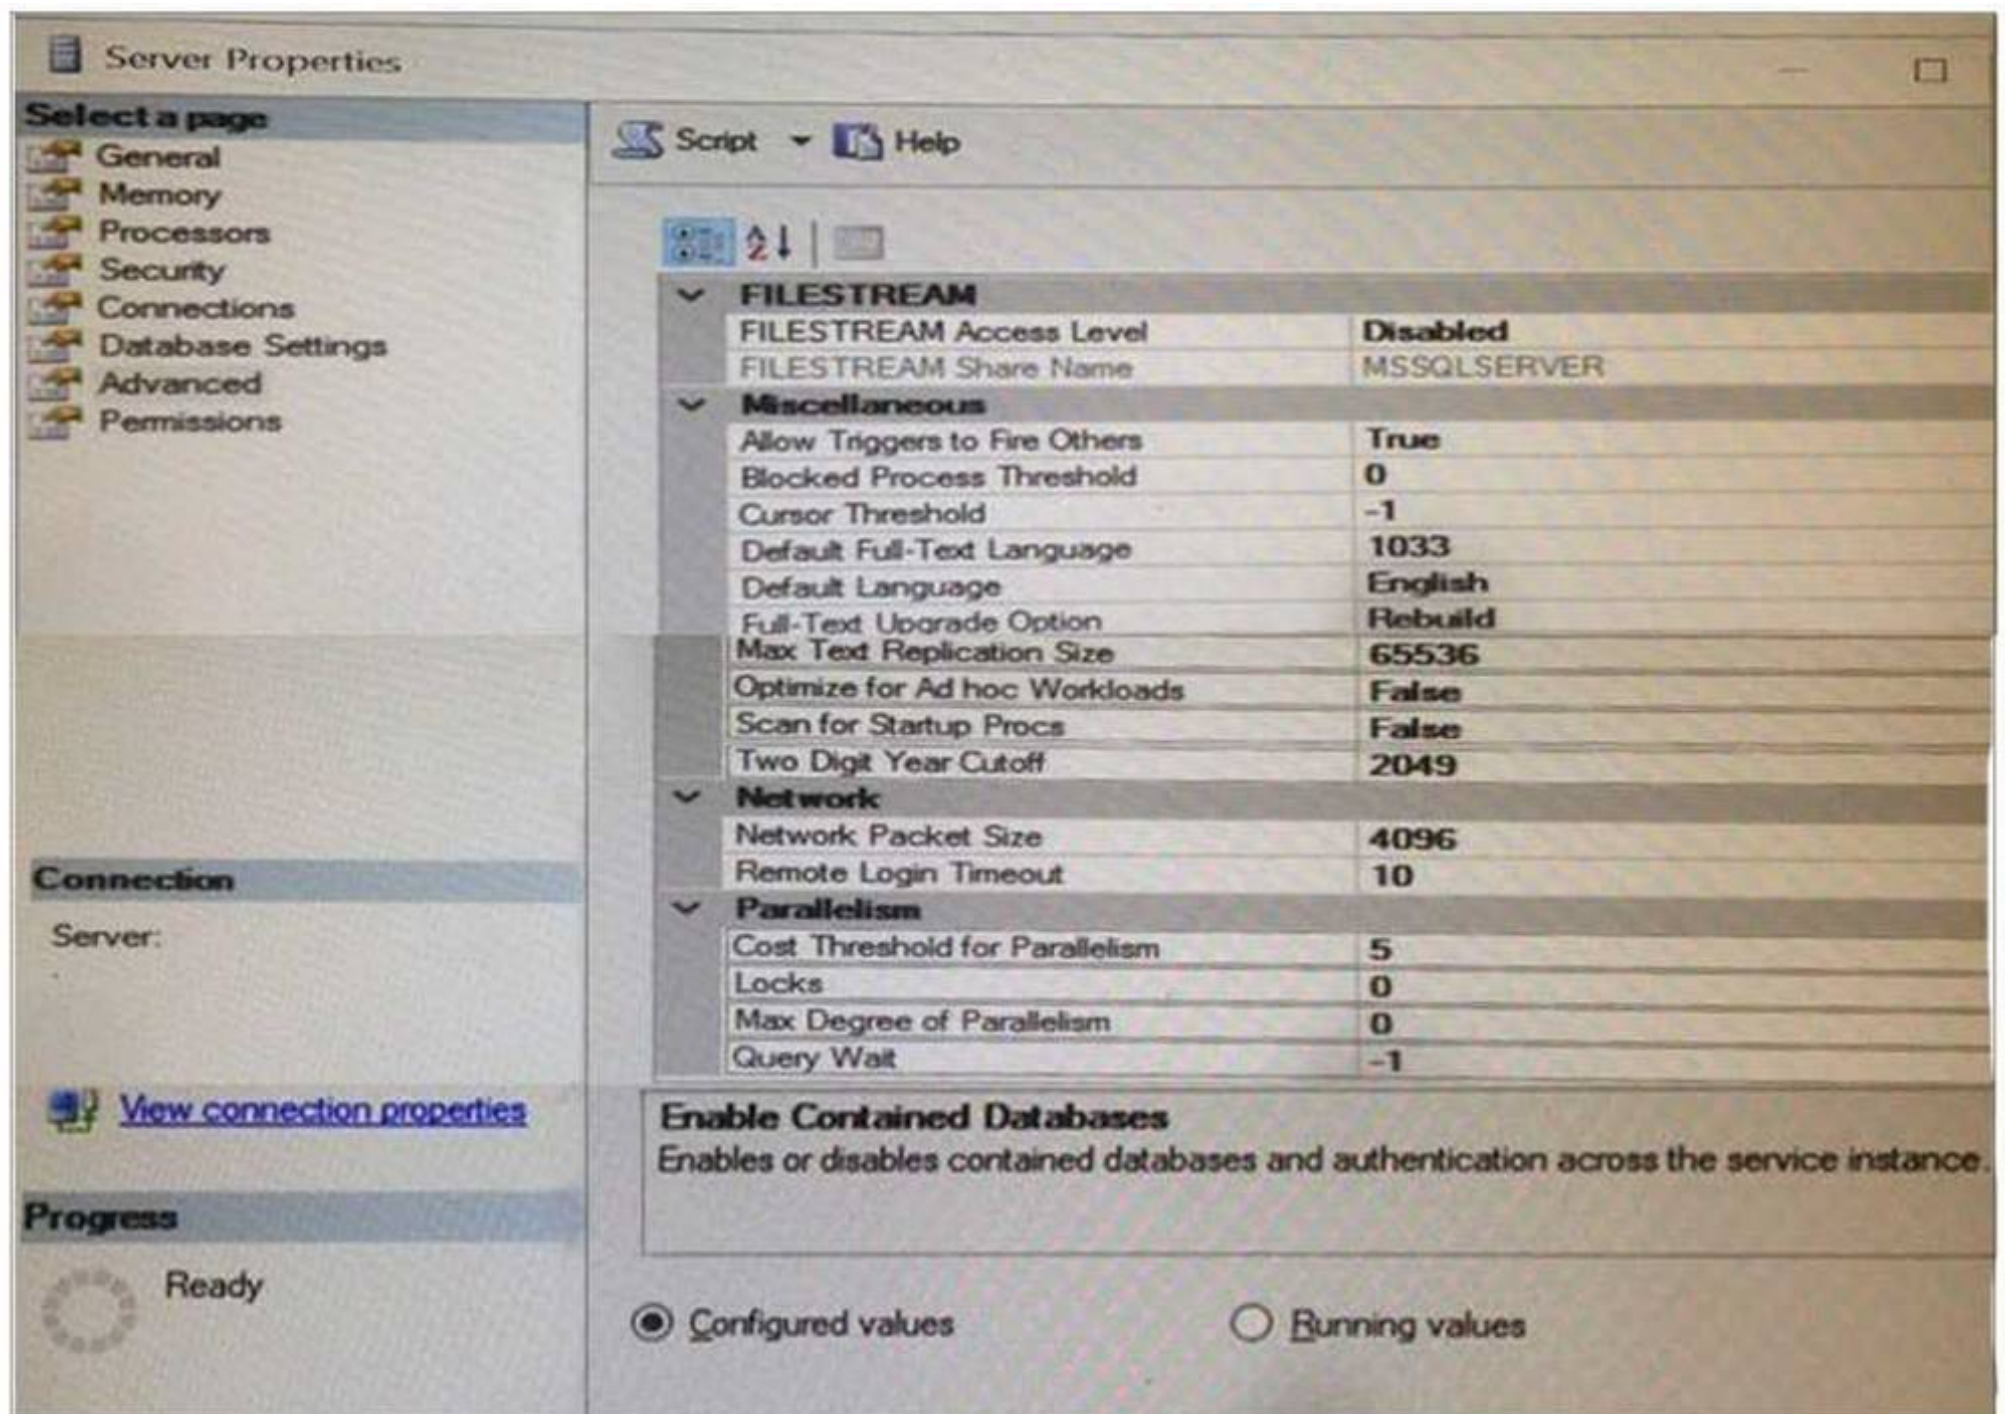
#### NEW QUESTION 26

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 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 27

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 32

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 question 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 increase the number of log files for the database. You store the log files across multiple disks. Does this meet the goal?

- A. Yes
- B. No

**Answer: A**

#### NEW QUESTION 34

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 the SQL Server instances in real time and optimize the server to maximize throughput, response time, and overall SQL performance.

You need to ensure that the performance of each instance is consistent for the same queried and query plans. 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: H**

**Explanation:** Advanced Viewing of Target Data from Extended Events in SQL Server

When your event session is currently active, you might want to watch the event data in real time, as it is received by the target.

Management > Extended Events > Sessions > [your-session] > Watch Live Data.

The query\_post\_execution\_showplan extended event enables you to see the actual query plan in the SQL Server Management Studio (SSMS) UI. When the Details pane is visible, you can see a graph of the query plan on the Query Plan tab. By hovering over a node on the query plan, you can see a list of property names and their values for the node.

The screenshot shows the SQL Server Management Studio (SSMS) interface. The 'Query Plan' tab is selected, displaying a query plan for 'Query 1: Query cost (relative to the batch): 100%'. The query plan consists of a 'Nested Loops (Left Outer Join)' operator, which is further broken down into 'Compute Scalar' and 'Nested Loops (Left Outer Join)' operators. A red box highlights the 'Query Plan' tab, and another red box highlights the 'Nested Loops (Left Outer Join)' operator. A red arrow points from the highlighted operator to the 'Properties' pane on the right, which displays detailed information about the 'Clustered Index Seek (Clustered)' operation.

**Clustered Index Seek (Clustered)**  
 Scanning a particular range of rows from a clustered index.

Physical Operation	Clustered Index Seek
Logical Operation	Clustered Index Seek
Actual Execution Mode	Row
Estimated Execution Mode	Row
Storage	RowStore
Actual Number of Rows	0
Actual Number of Batches	0
Estimated Operator Cost	0.0033263 (4%)
Estimated I/O Cost	0.003125
Estimated Subtree Cost	0.0033263
Estimated CPU Cost	0.0001581
Estimated Number of Executions	1.272901
Number of Executions	1
Estimated Number of Rows	1
Estimated Row Size	139 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	228

**Object**  
 [InMemTest2].[sys].[sysclsobjs].[clst] [s]

**Output List**  
 [InMemTest2].[sys].[sysclsobjs].name

**Seek Predicates**  
 Seek Keys[1]: Prefix: [InMemTest2].[sys].[sysclsobjs].class,  
 [InMemTest2].[sys].[sysclsobjs].id = Scalar Operator((50)),  
 Scalar Operator([InMemTest2].[sys].[sysnsobjs].[nsid] as [s],  
 [nsid])

References: <https://msdn.microsoft.com/en-us/library/mt752502.aspx>

#### NEW QUESTION 35

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 the SQL Server instances in real time and optimize the server to maximize throughput, response time, and overall SQL performance.

You need to collect query performance data while minimizing the performance impact on the SQL Server. 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 CollectorSet.
- 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:** C

**Explanation:** SQL Server Data Collector is a feature for performance monitoring and tuning available in SQL Server Management Studio.

Integration Services packages transform and load the collected data into the Microsoft Data Warehouse database.

Collection sets are defined and deployed on a server instance and can be run independently of each other. Each collection set can be applied to a target that matches the target types of all the collector types that are part of a collection set. The collection set is run by a SQL Server Agent job or jobs, and data is uploaded to the management data warehouse on a predefined schedule.

Predefined data collection sets include:

\* The Query Statistics data collection set collects information about query statistics, activity, execution plans and text on the SQL Server instance. It does not store all executed statements, only 10 worst performing ones.

\* Disk Usage data collection set collects information about disk space used by both data and log files for all databases on the SQL Server instance, growth trends, and average day growth. Etc.

References:

<http://www.sqlshack.com/sql-server-performance-monitoring-data-collector/>

#### NEW QUESTION 37

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.syssessions
- B. trace\_xenction\_map
- C. db
- D. syinotifications
- E. sysdbmaintplan.

**Answer:** A

#### NEW QUESTION 40

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 44

You are monitoring a Microsoft Azure SQL Database. The database is experiencing high CPU consumption.

You need to determine which query uses the most cumulative CPU.

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 one or not at all. You may need to drag the split bar between panes or scroll to view content.



#### Transact-SQL segments

#### Answer Area

sys.dm\_exec\_query\_stats o

sys.dm\_db\_partition\_stats o

sys.dm\_exec\_sessions o

sys.dm\_tran\_database\_transaction s o

highest\_cpu\_queries.plan\_handle DESC

highest\_cpu\_queries.total\_worker\_time DESC

q.objectid DESC

q.number DESC

```
SELECT
    highest_cpu_queries.plan_handle,
    highest_cpu_queries.total_worker_time,
    q.dbid,
    q.objectid,
    q.number,
    q.encrypted,
    q.[text]
FROM
    (SELECT TOP 50
        o.plan_handle,
        o.total_worker_time
    FROM
        Transact-SQL segment

    ORDER BY o.total_worker_time desc) AS highest_cpu_queries
CROSS APPLY sys.dm_exec_sql_text(plan_handle) AS q

ORDER BY Transact-SQL segment ;
```

#### Answer:

**Explanation:** Box 1: sys.dm\_exec\_query\_stats

sys.dm\_exec\_query\_stats returns aggregate performance statistics for cached query plans in SQL Server. Box 2: highest\_cpu\_queries.total\_worker\_time DESC  
 Sort on total\_worker\_time column

Example: The following example returns information about the top five queries ranked by average CPU time. This example aggregates the queries according to their query hash so that logically equivalent queries are grouped by their cumulative resource consumption.

USE AdventureWorks2012; GO

SELECT TOP 5 query\_stats.query\_hash AS "Query Hash",

SUM(query\_stats.total\_worker\_time) / SUM(query\_stats.execution\_count) AS "Avg CPU Time", MIN(query\_stats.statement\_text) AS "Statement Text"

FROM (SELECT QS.\*,

SUBSTRING(ST.text, (QS.statement\_start\_offset/2) + 1,

((CASE statement\_end\_offset

WHEN -1 THEN DATALENGTH(ST.text)

ELSE QS.statement\_end\_offset END

- QS.statement\_start\_offset)/2) + 1) AS statement\_text FROM sys.dm\_exec\_query\_stats AS QS

CROSS APPLY sys.dm\_exec\_sql\_text(QS.sql\_handle)as ST) as query\_stats GROUP BY query\_stats.query\_hash

ORDER BY 2 DESC;

References: <https://msdn.microsoft.com/en-us/library/ms189741.aspx>

#### NEW QUESTION 49

You have a Microsoft Azure SQL Database named MyDb that uses server version V12.

You plan to use Query Performance Insight to troubleshoot performance problems. The database query store is not enabled.

You need to enable the database query store to meet the following requirements for the database:

- Statistics must be aggregated every 15 minutes.
- Query stores must use no more than 1,024 megabytes (MB) of storage.
- Query information must be retained for at least 15 days.
- Queries must be captured based on resource consumption.

You connect to the database by using SQL Server Managements Studio.

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. Each correct selection is worth one point.

## Transact-SQL segments

STALE\_QUERY\_THRESHOLD\_DAYS

OFF

AUTO

ALL

DATA\_FLUSH\_INTERVAL\_SECONDS

INTERVAL\_LENGTH\_MINUTES

ON

## Answer area

ALTER DATABASE MyDb QUERY\_STORE = Transact-SQL segment

ALTER DATABASE MyDb SET QUERY\_STORE

```
(
    SIZE_BASED_CLEANUP_MODE = Transact-SQL segment ,
    CLEANUP_POLICY = ( Transact-SQL segment = 15),
    Transact-SQL segment = 15,
    QUERY_CAPTURE_MODE = Transact-SQL segment ,
    MAX_STORAGE_SIZE_MB = 1024
)
```

Answer:

Explanation:

## Transact-SQL segments

STALE\_QUERY\_THRESHOLD\_DAYS

OFF

AUTO

ALL

DATA\_FLUSH\_INTERVAL\_SECONDS

INTERVAL\_LENGTH\_MINUTES

ON

## Answer area

ALTER DATABASE MyDb QUERY\_STORE = ON

ALTER DATABASE MyDb SET QUERY\_STORE

```
(
    SIZE_BASED_CLEANUP_MODE = AUTO ,
    CLEANUP_POLICY = ( INTERVAL_LENGTH_MINUTES = 15),
    STALE_QUERY_THRESHOLD_DAYS = 15,
    QUERY_CAPTURE_MODE = AUTO ,
    MAX_STORAGE_SIZE_MB = 1024
)
```

or

## Transact-SQL segments

STALE\_QUERY\_THRESHOLD\_DAYS

OFF

AUTO

ALL

DATA\_FLUSH\_INTERVAL\_SECONDS

INTERVAL\_LENGTH\_MINUTES

ON

## Answer area

ALTER DATABASE MyDb QUERY\_STORE = ON

ALTER DATABASE MyDb SET QUERY\_STORE

```
(
    SIZE_BASED_CLEANUP_MODE = AUTO ,
    CLEANUP_POLICY = ( STALE_QUERY_THRESHOLD_DAYS = 15),
    INTERVAL_LENGTH_MINUTES = 15,
    QUERY_CAPTURE_MODE = AUTO ,
    MAX_STORAGE_SIZE_MB = 1024
)
```

Both answers are correct.

References: <https://msdn.microsoft.com/en-us/library/mt604821.aspx>

## NEW QUESTION 52

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(0) NOT NULL,
    DateCheclOut datetimeoffset(0) NOT NULL)
```

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:** Datetimeoffset defines a date that is combined with a time of a day that has time zone awareness and is based on a 24-hour clock.

Syntaxis: datetimeoffset [ (fractional seconds precision) ]

Forthe use"datetimeoffset(0)", the Fractional seconds precision is 0, which is required here. References: <https://msdn.microsoft.com/en-us/library/bb630289.aspx>

### NEW QUESTION 53

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.Female.Customers 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.

You need to create Sales.FemaleCustomers.

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.

#### Transact-SQL segments

#### Answer Area

WITH SCHEMABINDING

WITH ENCRYPTION

WITH CHECK OPTION

WITH VIEW\_METADATA

CREATE VIEW Website.Customer

Transact-SQL segments

```
AS SELECT s.CustomerNumber, s.CustomerName
FROM Sales.Customers as s
WHERE s.CustomerStatus=1
```

Transact-SQL segments

**Answer:**

**Explanation:** Box 1:WITH SCHEMABINDING:

SCHEMABINDING binds the view to the schema of the underlying table or tables. When SCHEMABINDING is specified, the base table or tables cannot be modified in a way that would affect the view definition.

Box 2: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: Sales.Female.Customers must meet the following requirements: References: <https://msdn.microsoft.com/en-us/library/ms187956.aspx>

#### NEW QUESTION 54

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.

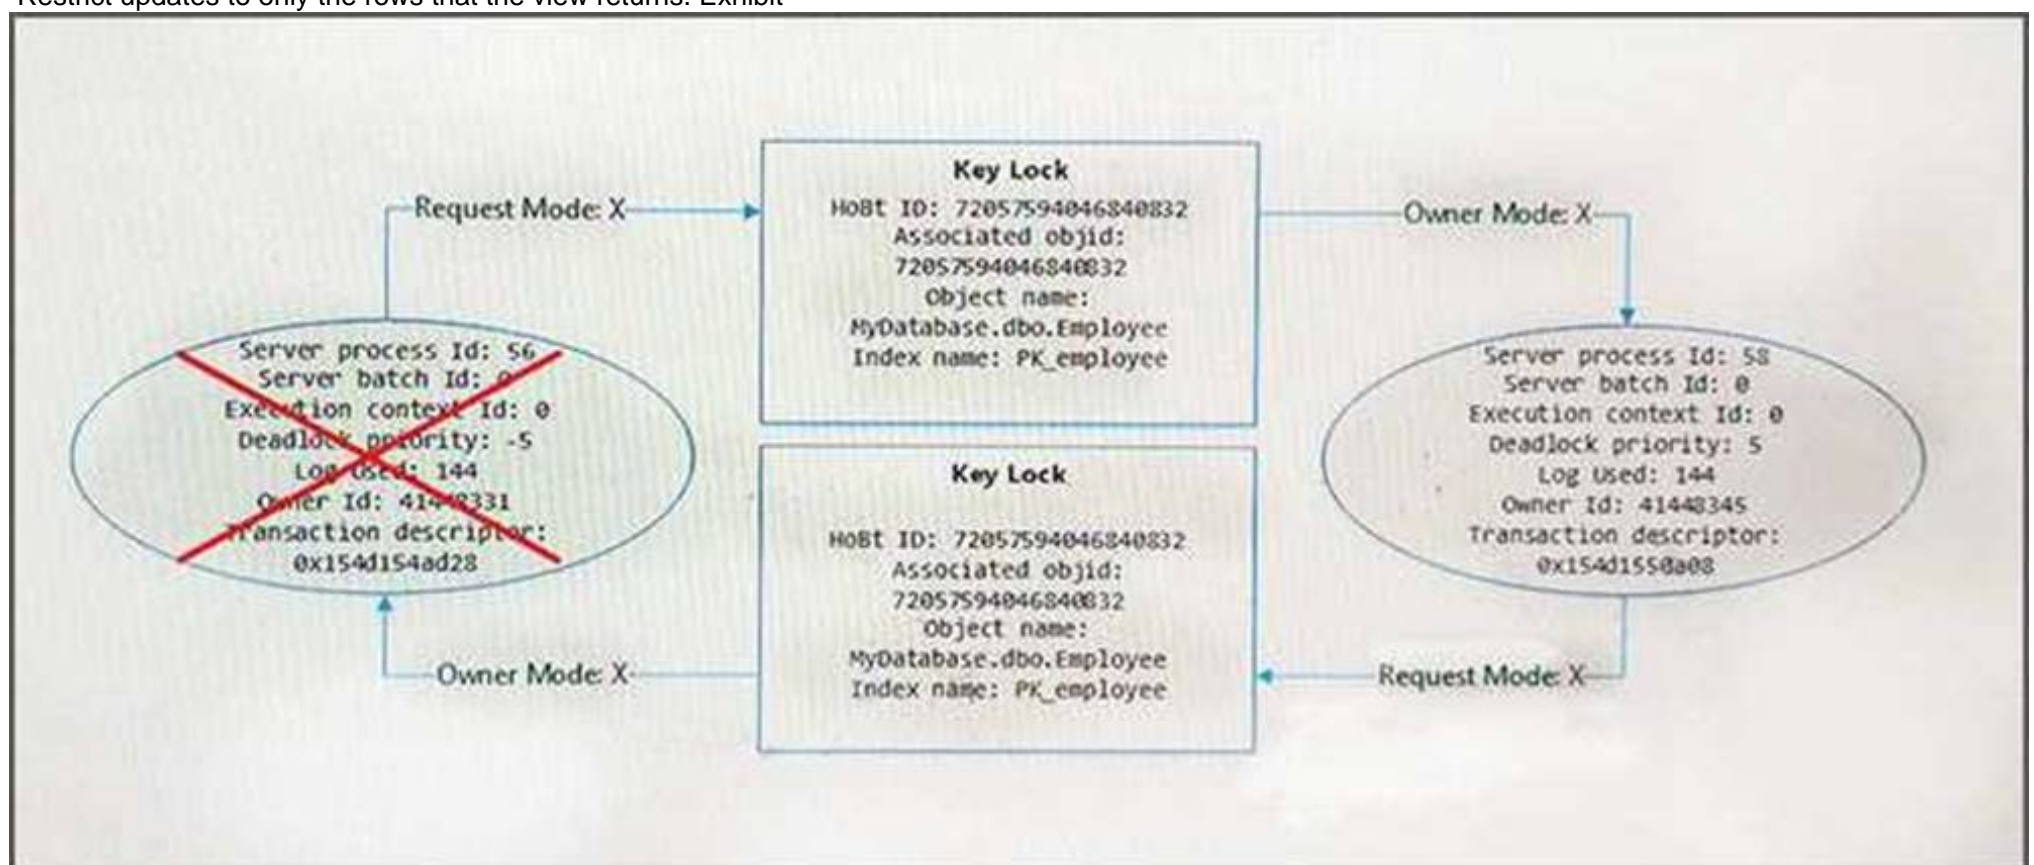
UspB: This stored procedure updates only the EmployeePayRate column.

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 are analyzing the performance of the database environment. You discover that locks that are held for a long period of time as the reports are generated. You need to generate the reports more quickly. The database must not use additional resources. What should you do?

- A. Update all FROM clauses of the DML statements to use the IGNORE\_CONSTRAINTS table hint.
- B. Modify the report queries to use the UNION statement to combine the results of two or more queries.
- C. Apply a nonclustered index to all tables used in the report queries.
- D. Update the transaction level of the report query session to READ UNCOMMITTED.

**Answer: D**

**Explanation:** Transactions running at the READ UNCOMMITTED level do not issue shared locks to prevent other transactions from modifying data read by the current transaction. This is the least restrictive of the isolation levels.

References: [https://technet.microsoft.com/en-us/library/ms173763\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms173763(v=sql.105).aspx)

#### NEW QUESTION 55

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 Microsoft SQL Server database named DB1 that contains the tables described in the following table:



Table name	Description
TBL1	<ul style="list-style-type: none"><li>• Column1 is configured as the primary key.</li><li>• The table will contain 20 million records.</li><li>• The table will contain historical data.</li><li>• Most queries of TBL1 return a high percentage of rows from the table with aggregates.</li></ul>
TBL2	<ul style="list-style-type: none"><li>• Column1 is configured as the primary key.</li><li>• The table will contain 25 million records.</li><li>• The frequency of updates and deletes to records in TBL2 is low.</li><li>• Most queries of TBL2 return a low percentage of rows and a high percentage of columns.</li></ul>

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

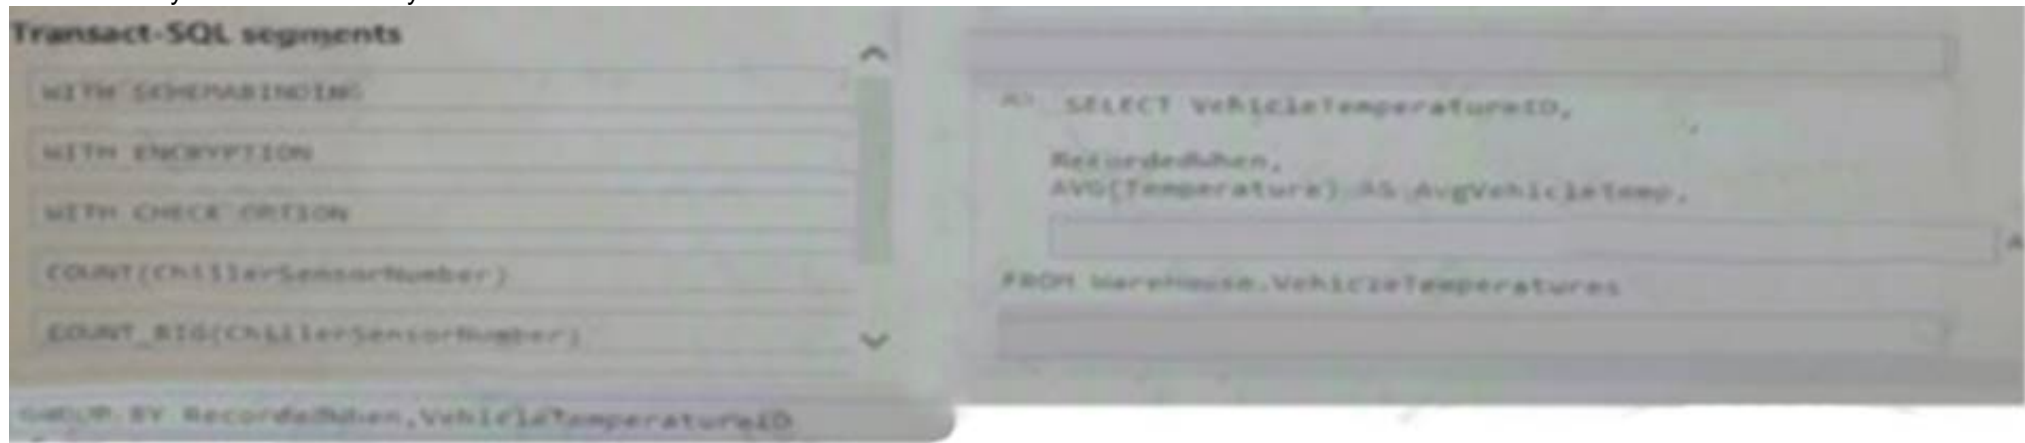
- A. Create clustered indexes on TBL1 and TBL2.
- B. Create a clustered index on TBL1. Create a nonclustered index on TBL2 and add the most frequently queried columns as included columns.
- C. Create a nonclustered index on tbl2 only.
- D. Create unique constraints on both TBL1 and TBL2.Create a partitioned 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 noncluster index on TBL2.
- G. Drop existing indexes on TBL1 and then create a clustered columnstore inde
- H. Create a nonclustered columnstore index on TBL1 Make no changes to TBL2.

Answer: B

NEW QUESTION 57

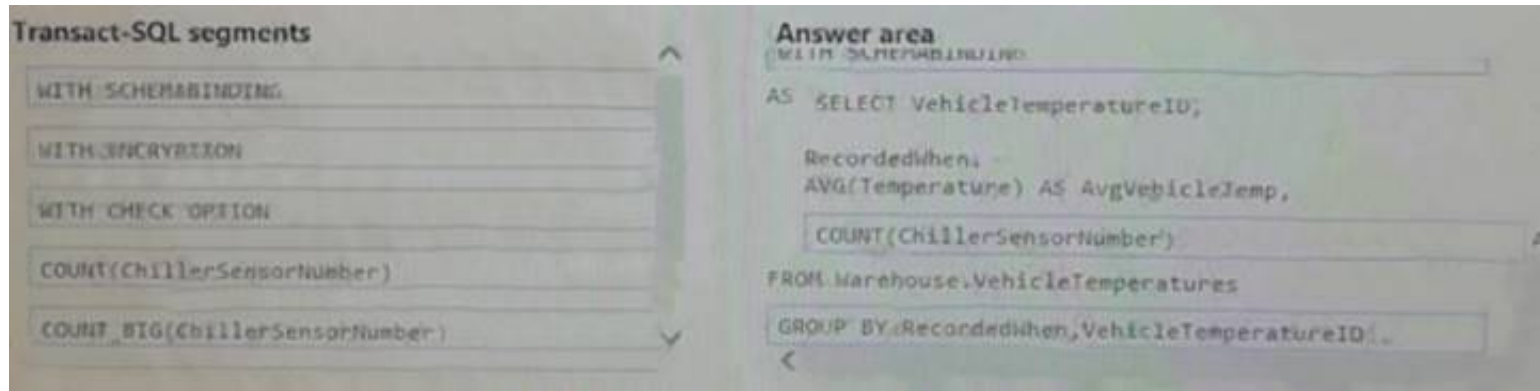
You are a database developer for a company that delivers produce and other refrigerated goods to grocery stores. You capture the food storage temperature for delivery vehicles by using Internet of Things (IoT) devices. You store the temperature data in a database table named vchicleTewperatures. The chillerSensorNumber column stores the identifier for the IoT devices.  
You need to create an indexed view that meets the following requirements:

- Persists the data on disk to reduce the amount of I/O.
- Provides the number of chillerSensorNumber items.
- Creates only a set sof summary rows.



Answer:

Explanation:



NEW QUESTION 58

You need to create a view that can be indexed. You write the following statement.

```
01 create view myView2
02
03 as
04 select col1, col2
05 from dbo.test
06 go
```

What should you add at line 02?

- A. with check\_option



- B. with recompile
- C. with view metadata
- D. with schenabinding

Answer: D

#### NEW QUESTION 59

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.

When a procedure calls spDeleteCustAcctRelationship, if the calling stored procedures has already started an active transaction, all the detections made by the spDeleteCustAccRelationship 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 soDeleteCustACCTRelationships 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 64

You are developing an app that allows users to query historical company financial data. You are reviewing email messages from the various stakeholders for a project.

The message from the security officer is shown in the Security Officer Email exhibit below. TO: Database developer

From: Security Officer

Subject: SQL object requirements

We need to simplify the security settings for the SQL objects. Having a assign permissions at every object in SQL is tedious and leads to a problem.

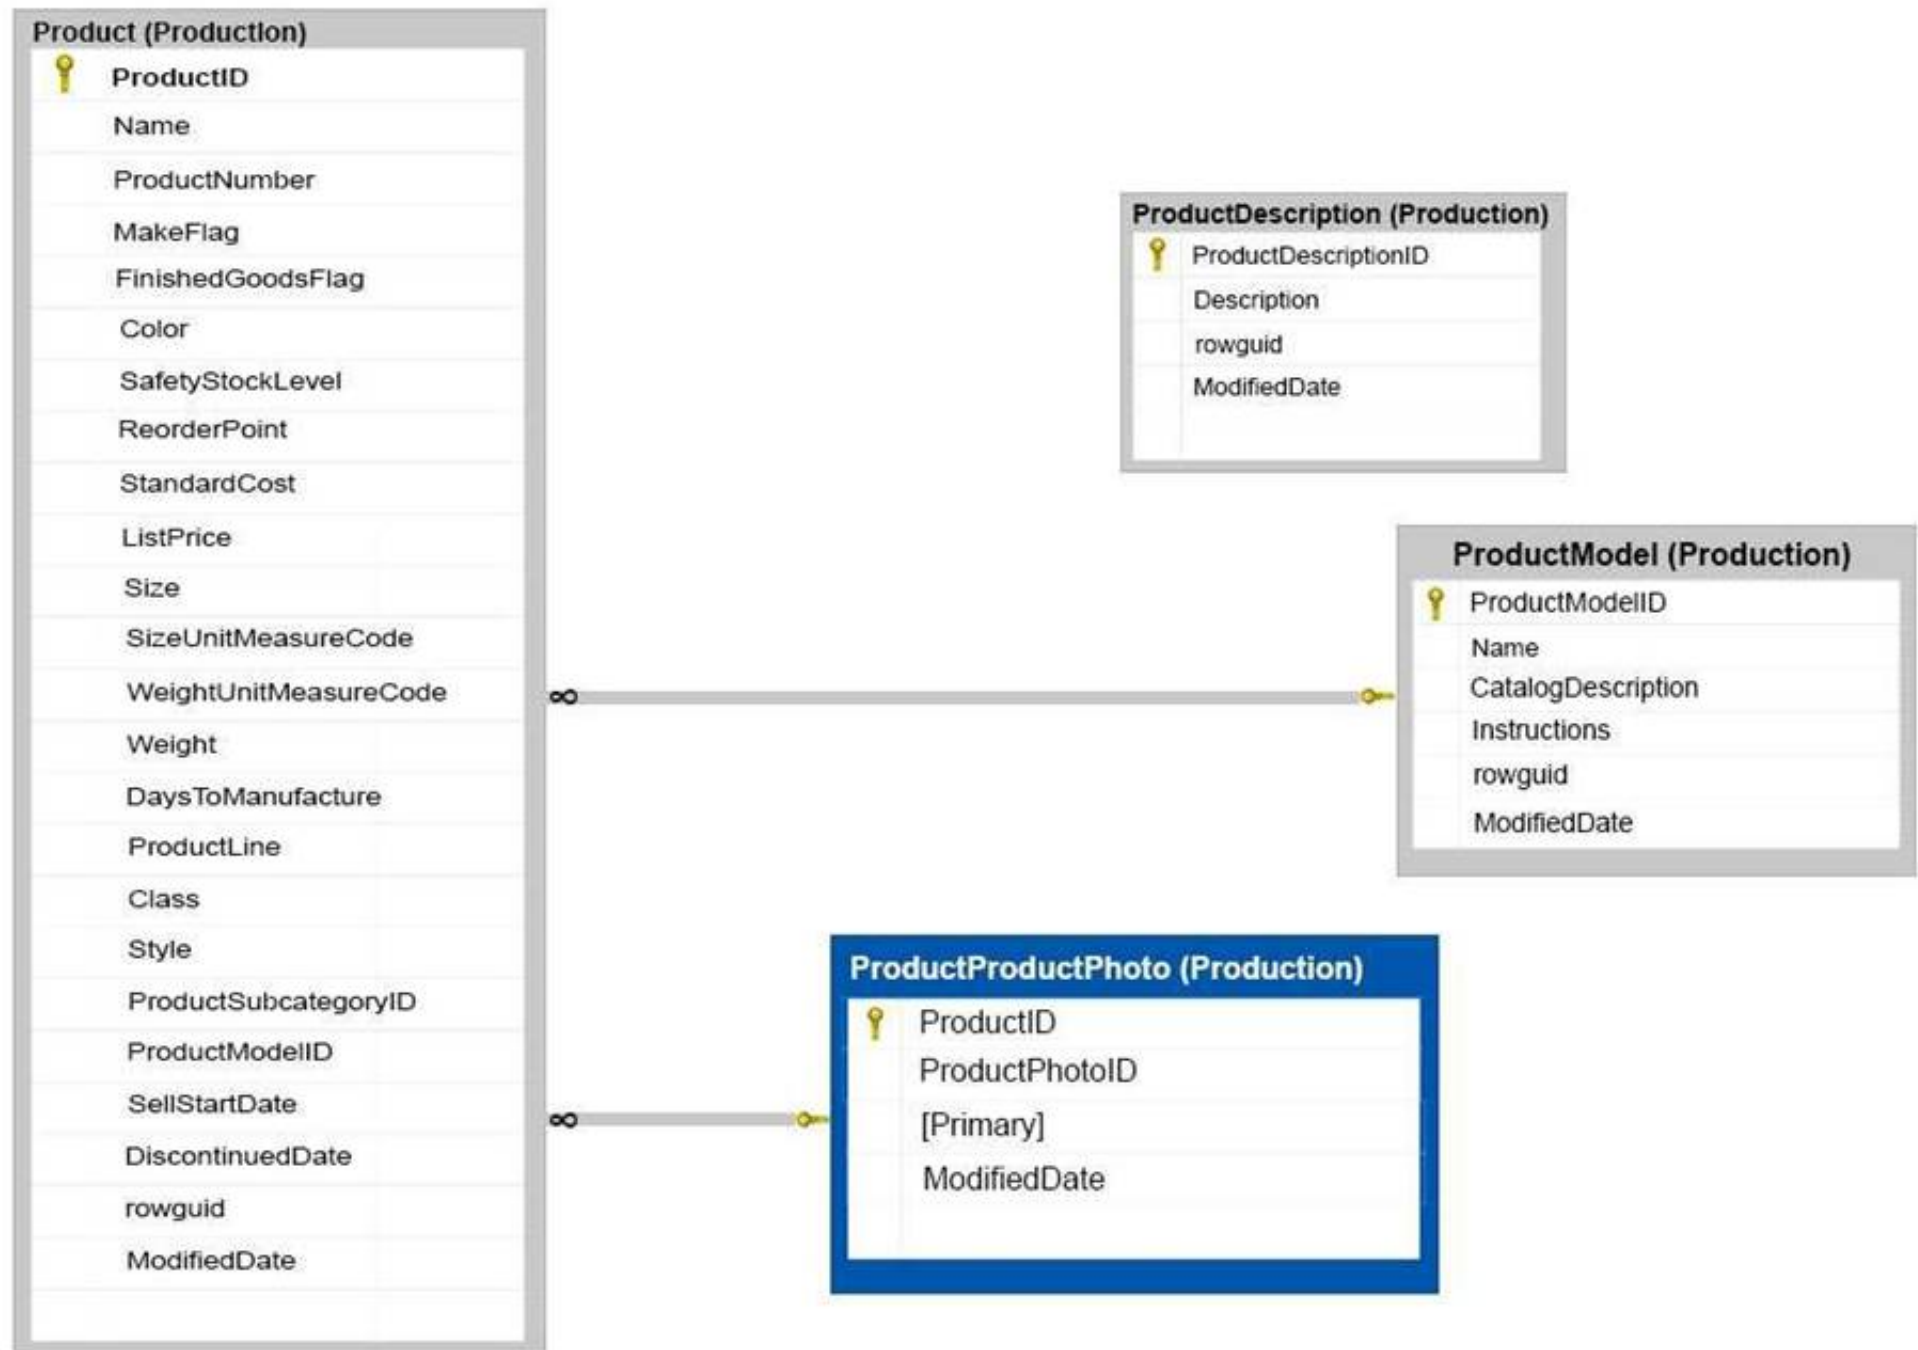
Documentation is also more difficult when we have to assign permissions at multiple levels. We need to assign the required permissions at one object, even though that object may be obtaining from other objects.

The message from the sales manager is shown in the Sales Manager Email exhibit below. TO: Database developer

From: Sales Manager Subject: Needed SQL objects

When creating objects for our use, they need to be flexible. We will be changing the base infrastructure frequently. We need components in SQL that will provide backward compatibility to our front end applications as the environments change so that do not need to modify the front end applications. We need objects that can provide a filtered set of the data. The data may be coming from multiple tables and we need an object that can provide access to all of the data through a single object reference.

This is an example of the types of data we need to be able to have queries against without having to change the front end applications.



The message from the web developer is shown in the Web Developer Email exhibit below. TO: Database developer

From: Web Developer

Subject: SQL Object component

Whatever you will be configuring to provide access to data in SQL, it needs to connect using the items referenced in this interface. We have been using this for a long time, and we cannot change this from end easily. Whatever objects are going to be used in SQL they must work using object types this interface references.

Database Name:

Table Name:

Column Name:

You need to create one or more objects that meet the needs of the security officer, the sales manager and the web developer. For each of the following statements, select Yes if the statement is true. Otherwise, select No.



## Answer Area

	Yes	No
You must create a stored procedure	<input type="radio"/>	<input type="radio"/>
You must create a trigger	<input type="radio"/>	<input type="radio"/>
You must create a view	<input type="radio"/>	<input type="radio"/>

### Answer:

**Explanation:** \* Stored procedure: Yes

A stored procedure to implement the following:

Whatever you will be configuring to provide access to data in SQL, it needs to connect using the items referenced in this interface. We have been using this for a long time, and we cannot change this from end easily. Whatever objects are going to be used in SQL they must work using object types this interface references.

\* Trigger: No

No requirements are related to actions taken when changing the data.

\* View: Yes

Because: We need objects that can provide a filtered set of the data. The data may be coming from multiple tables and we need an object that can provide access to all of the data through a single object reference.

### NEW QUESTION 65

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 70

You are creating the following two stored procedures:  
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 71

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 73

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 the SQL Server instances in real time and optimize the server to maximize throughput, response time, and overall SQL performance.

What should you do?

- A. Create asys.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:** B

**Explanation:** 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. This information includes client version, client program name, client login time, login user, current session setting, and more. Use sys.dm\_exec\_sessions to first view the current system load and to identify a session of interest, and then learn more information about that session by using other dynamic management views or dynamic management functions.

Examples of use include finding long-running cursors, and finding idle sessions that have open transactions.

### NEW QUESTION 75

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	decimail(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?

- Create a clustered index on the table.
- Create a nonclustered index on the table.
- Create a nonclustered filtered index on the table.
- Create a clustered columnstore index on the table.
- Create a nonclustered columnstore index on the table.
- 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 80

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 have a table and a stored procedure that were created by running the following Transact-SQL statements:

```
CREATE TABLE Employee
(
    EmployeeId int NOT NULL PRIMARY KEY,
    FirstName varchar(20),
    LastName varchar(20),
    Status char(1),
    Address varchar(100),
    Department int NOT NULL
);

CREATE PROCEDURE uspSelectEmployeeDetails
(
    @LastName varchar(20)
)
AS
BEGIN
    SELECT e.FirstName, e.LastName, d.DepartmentName
    FROM Employee e
    JOIN Department d on e.DepartmentId = d.DepartmentId
    WHERE e.Status = 'T' AND e.LastName = @LastName;
END;
```

The Employee table is persisted on disk. You add 2,000 records to the Employee table. You need to create an index that meets the following requirements:

- Optimizes the performance of the stored procedure.
- Covers all the columns required from the Employee table.
- Uses FirstName and LastName as included columns.
- Minimizes index storage size and index key size. What should you do?

- Create a clustered index on the table.
- Create a nonclustered index on the table.
- Create a nonclustered filtered index on the table.
- Create a clustered columnstore index on the table.
- Create a nonclustered columnstore index on the table.

F. Create a hash index on the table.

**Answer:** B

**Explanation:** References: [https://technet.microsoft.com/en-us/library/jj835095\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/jj835095(v=sql.110).aspx)

#### NEW QUESTION 84

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 85

You plan to create a stored procedure that uses a table parameter as an input parameter. The table value parameter may hold between 1 and 10,000 rows when you run the stored procedure.

The stored procedure will use the rows within the table value parameter to filter the rows that will be returned by the SELECT statement.

You need to create the stored procedure and ensure that it runs as quickly as possible.

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

NOTE: Each correct answer selection is worth one point.

## Answer Area

CREATE PROCEDURE GetRows

▼	MyDataType
MyInputTable@ MyInputTable# MyDataType@ MyDataType#	

AS

▼	(Ids INT PRIMARY KEY)
CREATE TABLE #MyInputTable DECLARE @MyInputTable TABLE DECLARE @MyInputTable TABLE (MyDataType int NOT NULL	

INSERT INTO

SELECT Ids

FROM @MyInputTable

▼
MyInputTable# MyInputTable@ MyDatatype@

\*.SELECT t

FROM FACT\_Customers t

▼
JOIN #MyInputTable m JOIN @MyInputTable m JOIN @MyDataType m

ON t.CustomerId = m.Ids

**Answer:**

**Explanation:**

## Answer Area

```
CREATE PROCEDURE GetRows
```

```
    MyInputTable@  
    MyInputTable#  
    MyDataType@  
    MyDataType#
```

```
AS
```

```
    CREATE TABLE #MyInputTable  
    DECLARE @MyInputTable TABLE  
    DECLARE @MyInputTable TABLE (MyDataType int NOT NULL
```

```
INSERT INTO
```

```
SELECT Ids
```

```
FROM @MyInputTable
```

```
*.SELECT t
```

```
FROM FACT_Customers t
```

```
    JOIN #MyInputTable m  
    JOIN @MyInputTable m  
    JOIN @MyDataType m
```

```
ON t.CustomerId = m.Ids
```

### NEW QUESTION 88

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 table
- B. Create a natively compiled stored procedure to cleanse and consolidate the data.
- C. Convert the reference table to a memory-optimized table
- D. Set the durability option to SCHEMA\_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 table
- G. Set the durability option to SCHEMA\_ONLY.

**Answer:** D

### NEW QUESTION 91

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 defragment both indexes. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

**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 92

You are evaluating the performance of a database environment.

You must avoid unnecessary locks and ensure that lost updates do not occur. You need to choose the transaction isolation level for each data scenario.

Which isolation level should you use for each scenario? To answer, drag the appropriate isolation levels to the correct scenarios. Each isolation 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

### Answer area

- read committed
- serializable
- read uncommitted
- repeatable read

### Scenario

Reading accurate data is top priority. Select statements will wait until any transaction that currently owns the data has been committed or rolled back before returning the value

Performance is top priority. The work and memory required by the Microsoft SQL Server lock manager is reduced

The same select statement is issued multiple times within a transaction and the same result are returned. New records are allowed to be inserted into the table referenced by the Select statement

### Isolation levels

Isolation level

Isolation level

Isolation level

Answer:

**Explanation:** Box 1: Readcommitted

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. Box 2: Read Uncommitted

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. Box 3: Serializable

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 94

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 the SQL Server instances in real time and optimize the server to maximize throughput, response time, and overall SQL performance.

You need to create a baseline set of metrics to report how the computer running SQL Server operates under normal load. The baseline must include the resource usage associated with the server processes. 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 asys.dm\_os\_wait\_stats query.
- H. Create an Extended Event.

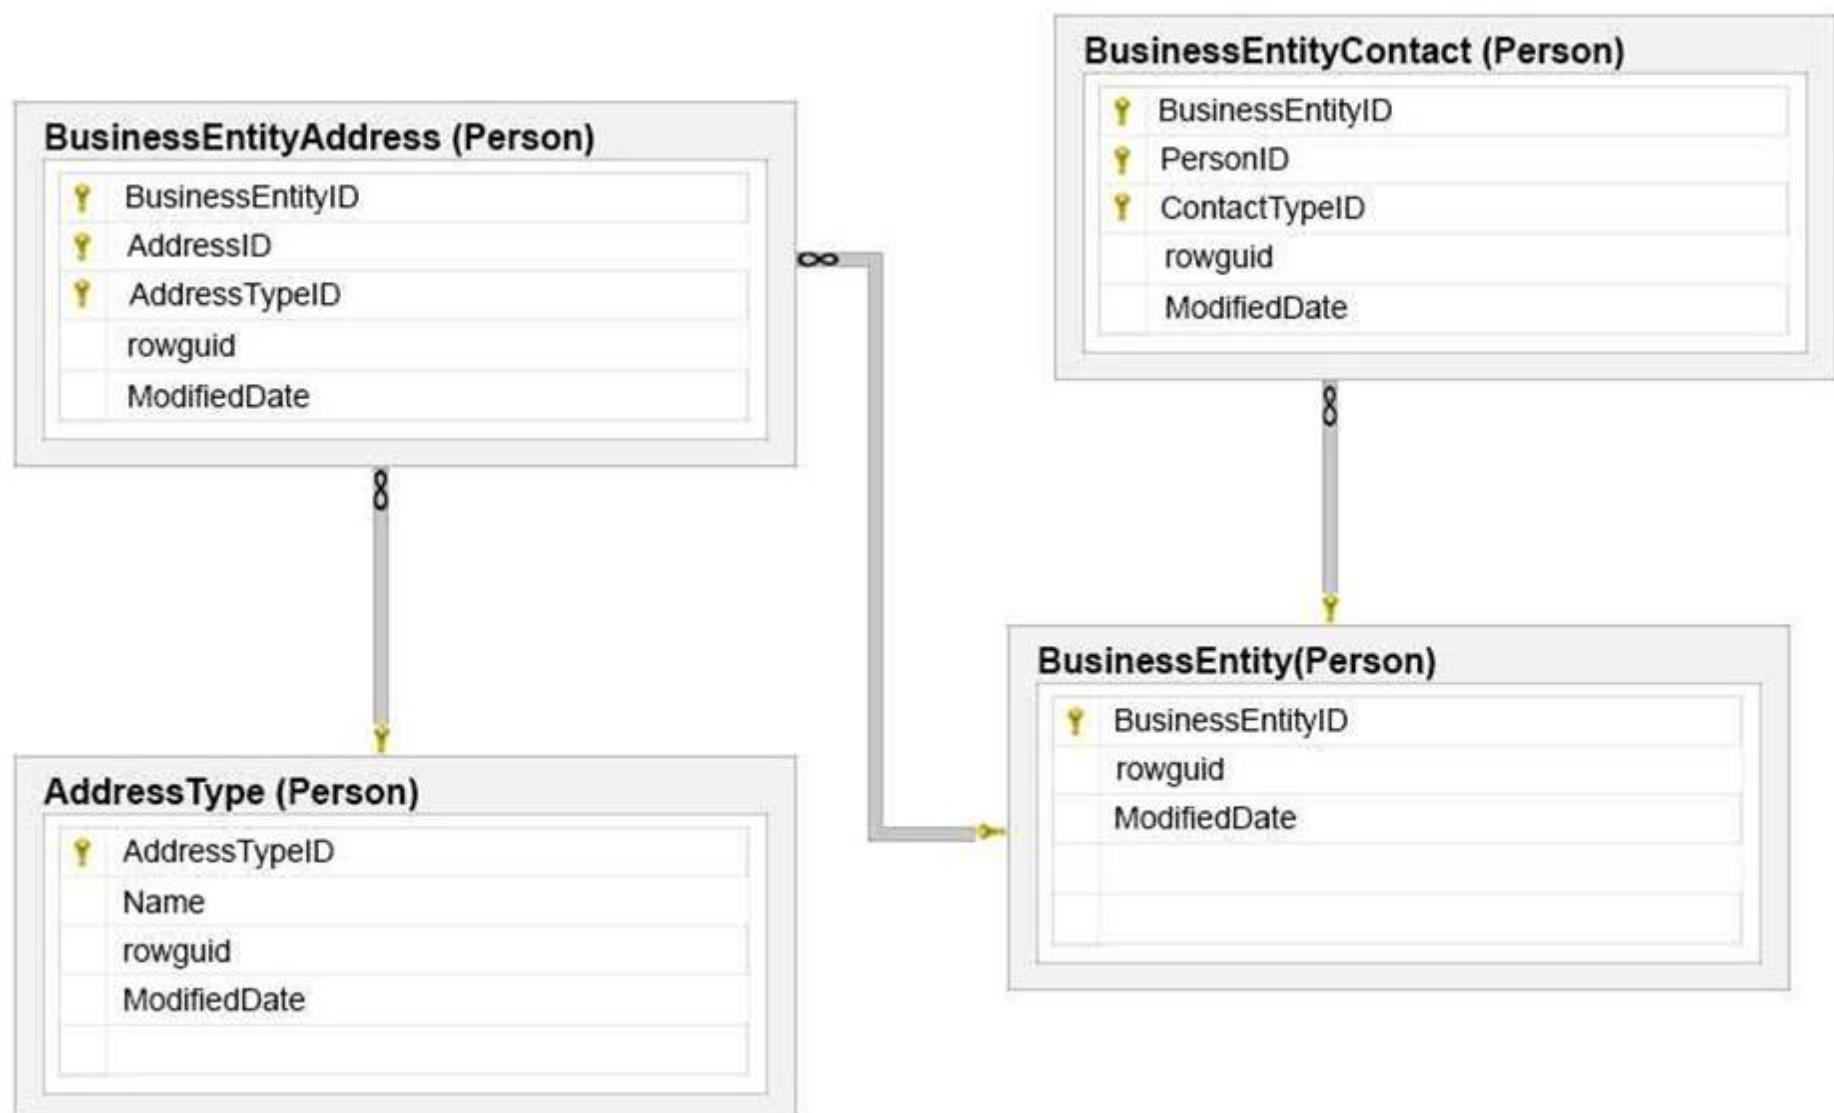
**Answer: D**

**Explanation:** sys.dm\_os\_memory\_objects returns memory objects that are currently allocated by SQL Server. You can use sys.dm\_os\_memory\_objects to analyze memory use and to identify possible memory leaks.

Example: The following example returns the amount of memory allocated by each memory object type. SELECT SUM (pages\_in\_bytes) as 'Bytes Used', type FROM sys.dm\_os\_memory\_objects GROUP BY type ORDER BY 'Bytes Used' DESC; GO

#### NEW QUESTION 98

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 the, 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 102

You manage a Microsoft Azure SQL Database that has the Standard tier plan. The database size has increased and users experience slow performance. You need to identify usage for the following resources:

- CPU utilization
- Disk storage
- Memory utilization
- Disk I/O

.....

- A. Activity monitor
- B. sys.dm\_exec\_connections
- C. Azure portal
- D. sys.rtuource\_usage
- E. sys . dm\_db\_resource\_stats

**Answer:** BC

### NEW QUESTION 103

Case Study 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
```

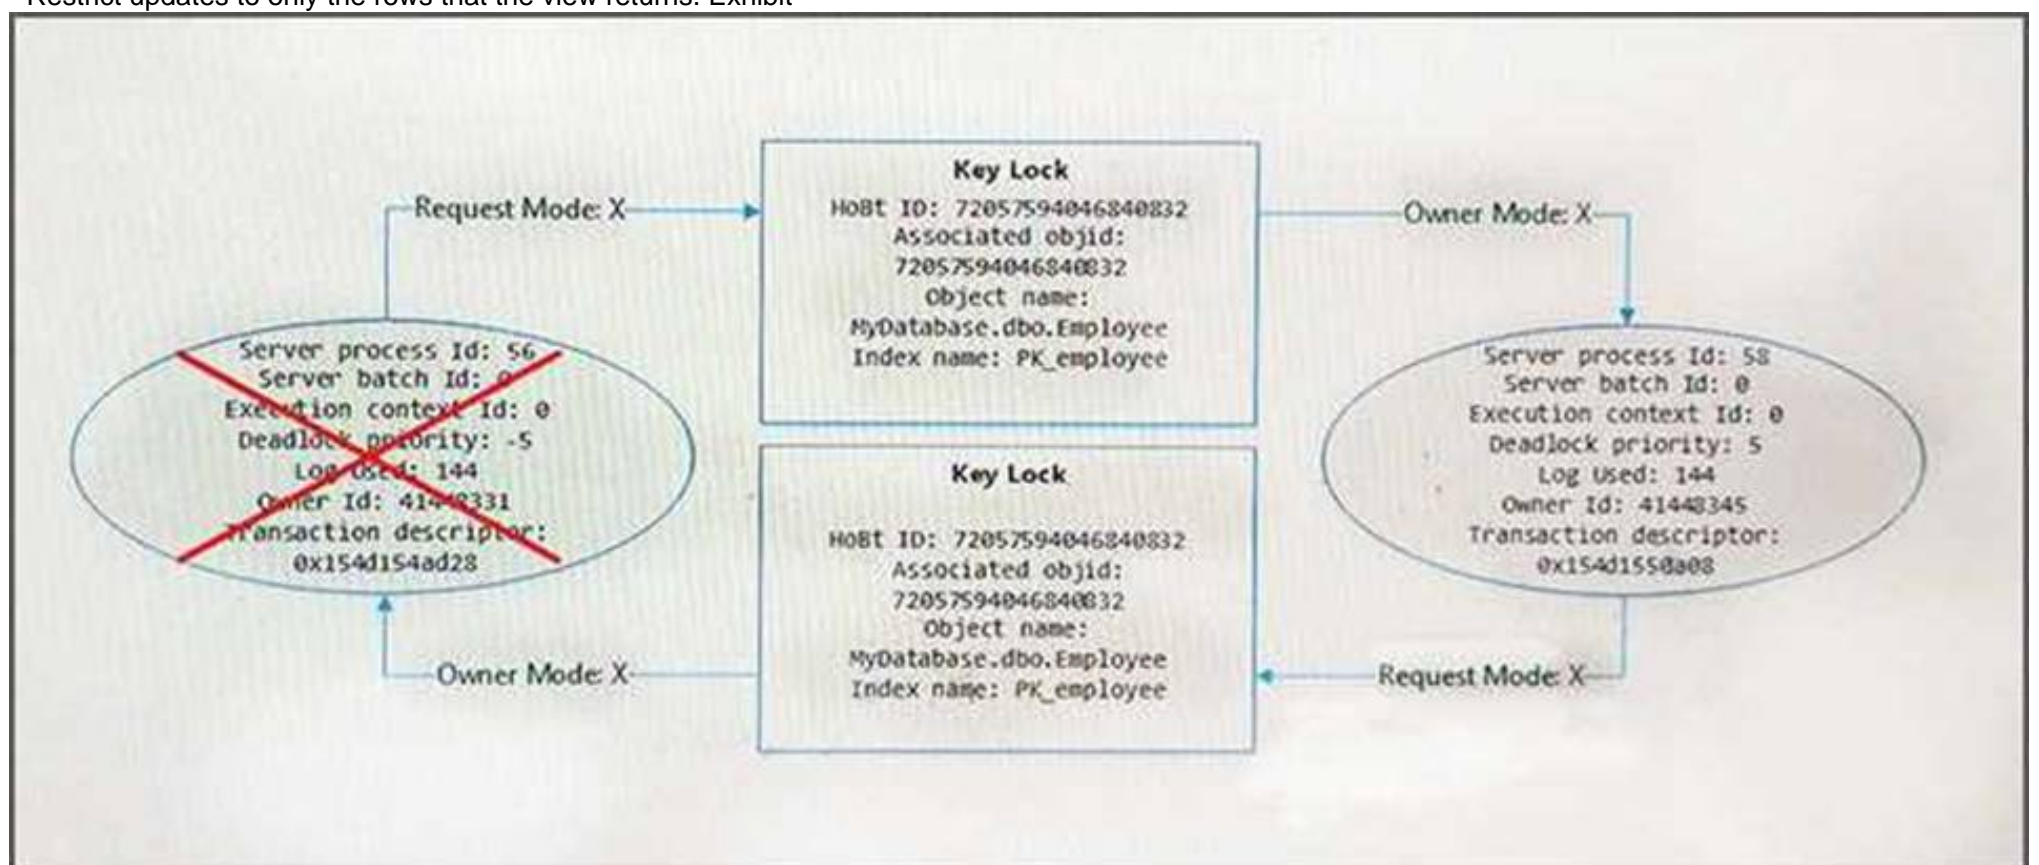
#### Application

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



Both of the stored procedures experience blocking issues. UspB must not abort if UspA commits changes to a row before UspB commits changes to the same row. UspA must not abort if UspB commits changes to a row before UspA commits changes to the same row.

You need to specify the transaction isolation levels to enable row versioning.

How should you complete the Transact-SQL statements? To answer, drag the 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

SET ALLOW\_SNAPSHOT\_ISOLATION OFF;

SET TRANSACTION ISOLATION LEVEL  
READ COMMITTED;

SET TRANSACTION ISOLATION LEVEL  
READ UNCOMMITTED;

SET TRANSACTION ISOLATION LEVEL  
REPEATABLE READ;

SET TRANSACTION ISOLATION LEVEL  
SERIALIZABLE;

SET READ\_COMMITTED\_SNAPSHOT ON;

SET READ\_COMMITTED\_SNAPSHOT OFF;

## Answer area

ALTER DATABASE HR1

Transact-SQL segment

CREATE PROCEDURE UspA  
AS  
BEGIN

SET NOCOUNT ON;

Transact-SQL segment

BEGIN TRANSACTION;

...

COMMIT TRANSACTION;

END;  
CREATE PROCEDURE UspB  
AS  
BEGIN

Transact-SQL segment

BEGIN TRANSACTION;

...

COMMIT TRANSACTION;

Answer:

**Explanation:** References: [https://technet.microsoft.com/en-us/library/ms175095\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms175095(v=sql.105).aspx) [https://technet.microsoft.com/en-us/library/ms173763\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms173763(v=sql.105).aspx)

### NEW QUESTION 108

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

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><div></div><div>interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function</div></div>
Module2	<div><div></div><div>interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function</div></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 110

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 112

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 116

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

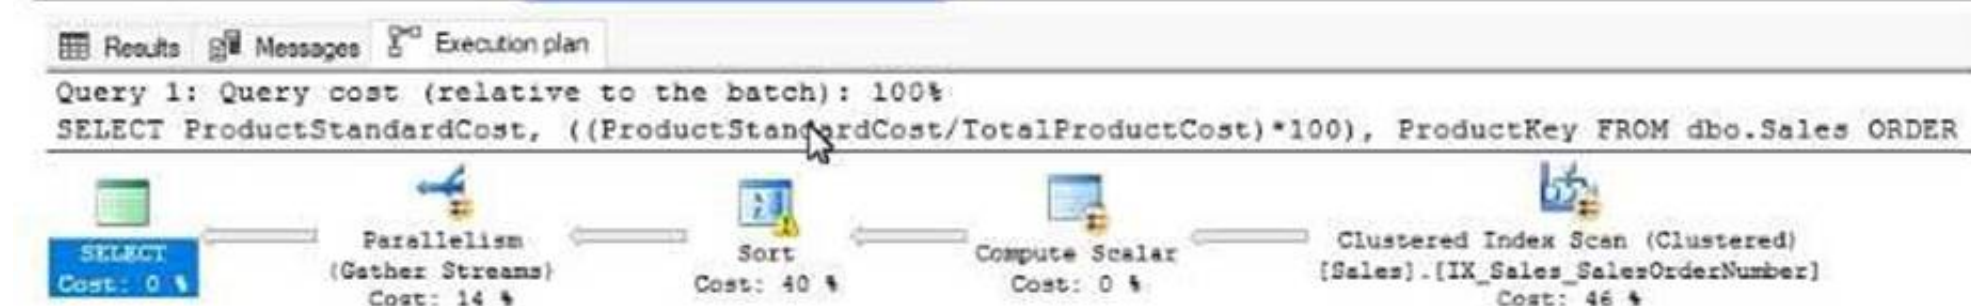
### NEW QUESTION 120

You have a Microsoft SQL Server database that has a table named Sales. The table is used for retrieving data and is updated during non business hours.

You run the following Transact-SQL statement:

```
SELECT ProductStandardCost, ((ProductStandardCost/TotalProductCost)*100) StandardCostRatio, ProductKey
FROM Sales
ORDER BY ProductStandardCost DESC, ProductKey
```

You analyze the execution plan for the statement. (Click the Exhibit button).



You need to add an index that optimizes performance. How should you complete the Transact-SQL statement?

```
CREATE INDEX IX_Sales_01
ON Sales (
    ProductStandardCost DESC
    WITH (ONLINE = ON)
    INCLUDE (TotalProductCost)
    WITH (SORT_IN_TEMPDB = ON)
    WITH (SORT_IN_TEMPDB = OFF)
)
```

Answer:

Explanation:



### NEW QUESTION 121

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 Standard service tier. Does this meet the goal?

- A. Yes
- B. No



**Answer:** B

**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 124

You suspect deadlocks on a database.

Which two trace flags in the Microsoft SQL Server error log should you locate? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. 1204
- B. 1211
- C. 1222
- D. 2528
- E. 3205

**Answer:** AC

**Explanation:** Trace flag 1204 returns the resources and types of locks participating in a deadlock and also the current command affected.

Trace flag 1222 returns the resources and types of locks that are participating in a deadlock and also the current command affected, in an XML format that does not comply with any XSD schema.

References:

[https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-traceon-trace-flags-transact-sql?vie](https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-traceon-trace-flags-transact-sql?view=sql-server-11)

#### NEW QUESTION 127

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

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 129**

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 stored 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 stored procedure on the Employees table. Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:** We should use table-valued functions, not procedures, to customize the recorded change data. References: <https://msdn.microsoft.com/en-us/library/cc645858.aspx>

**NEW QUESTION 133**

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
```

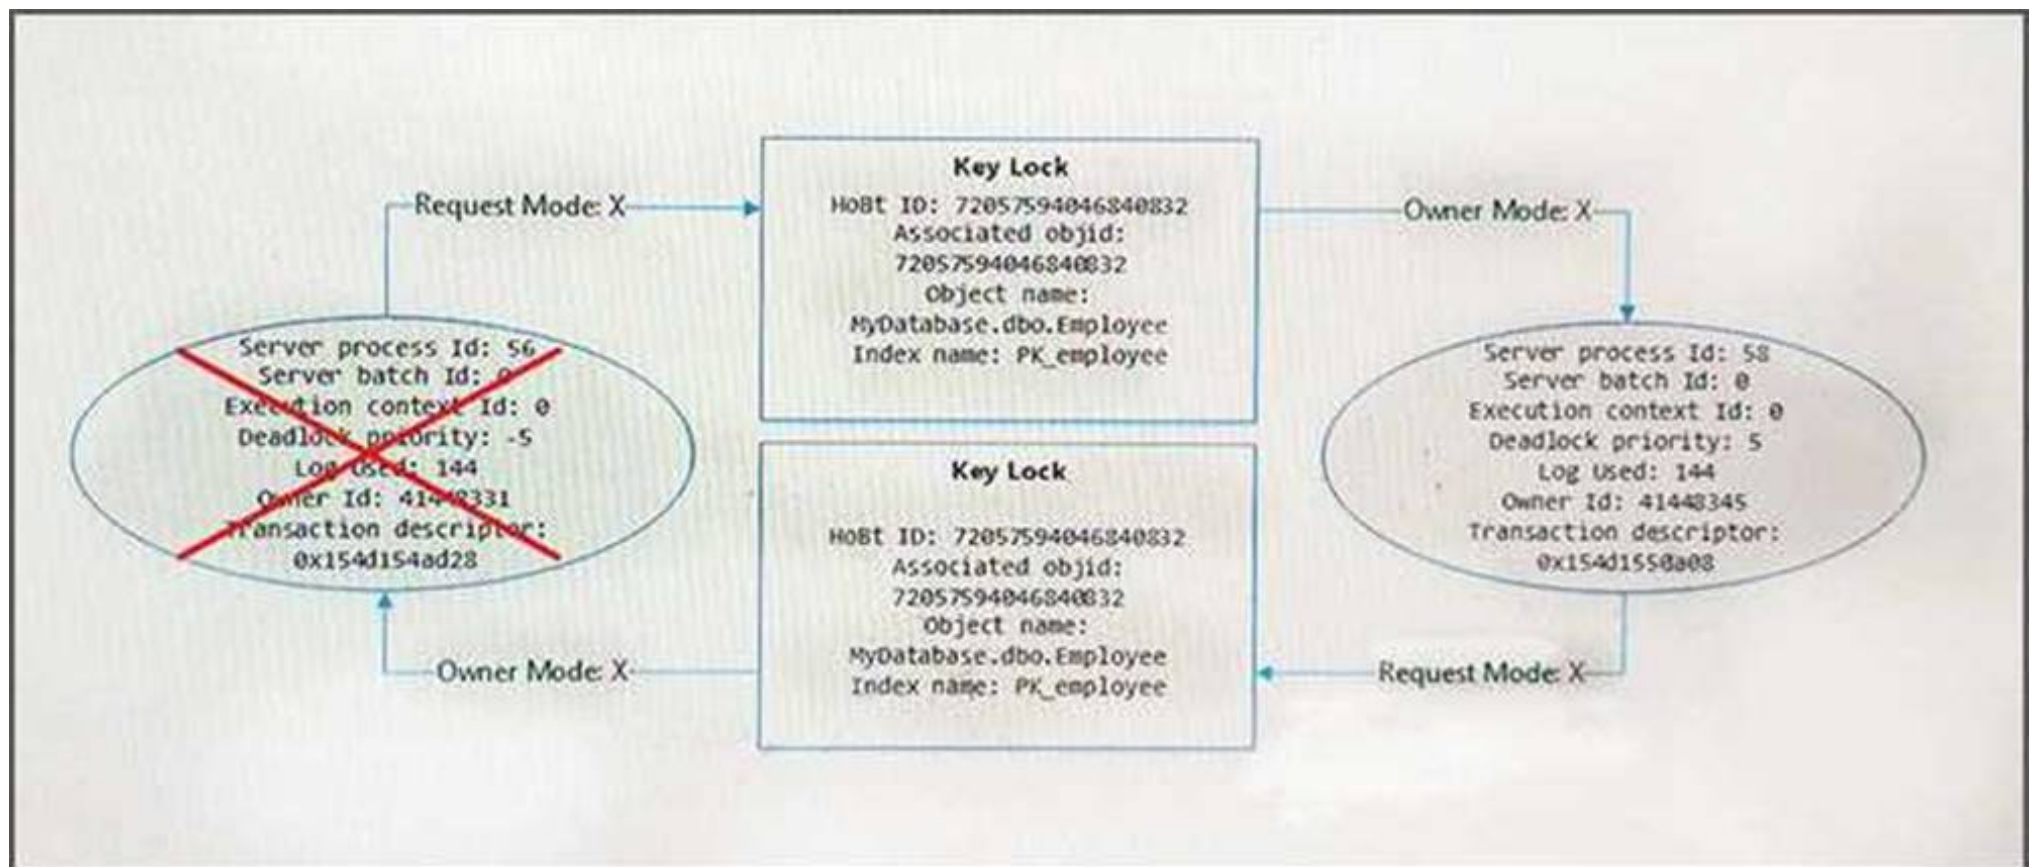
Application

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 136

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 137

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 check constraint on the table.  
Does the solution meet the goal?

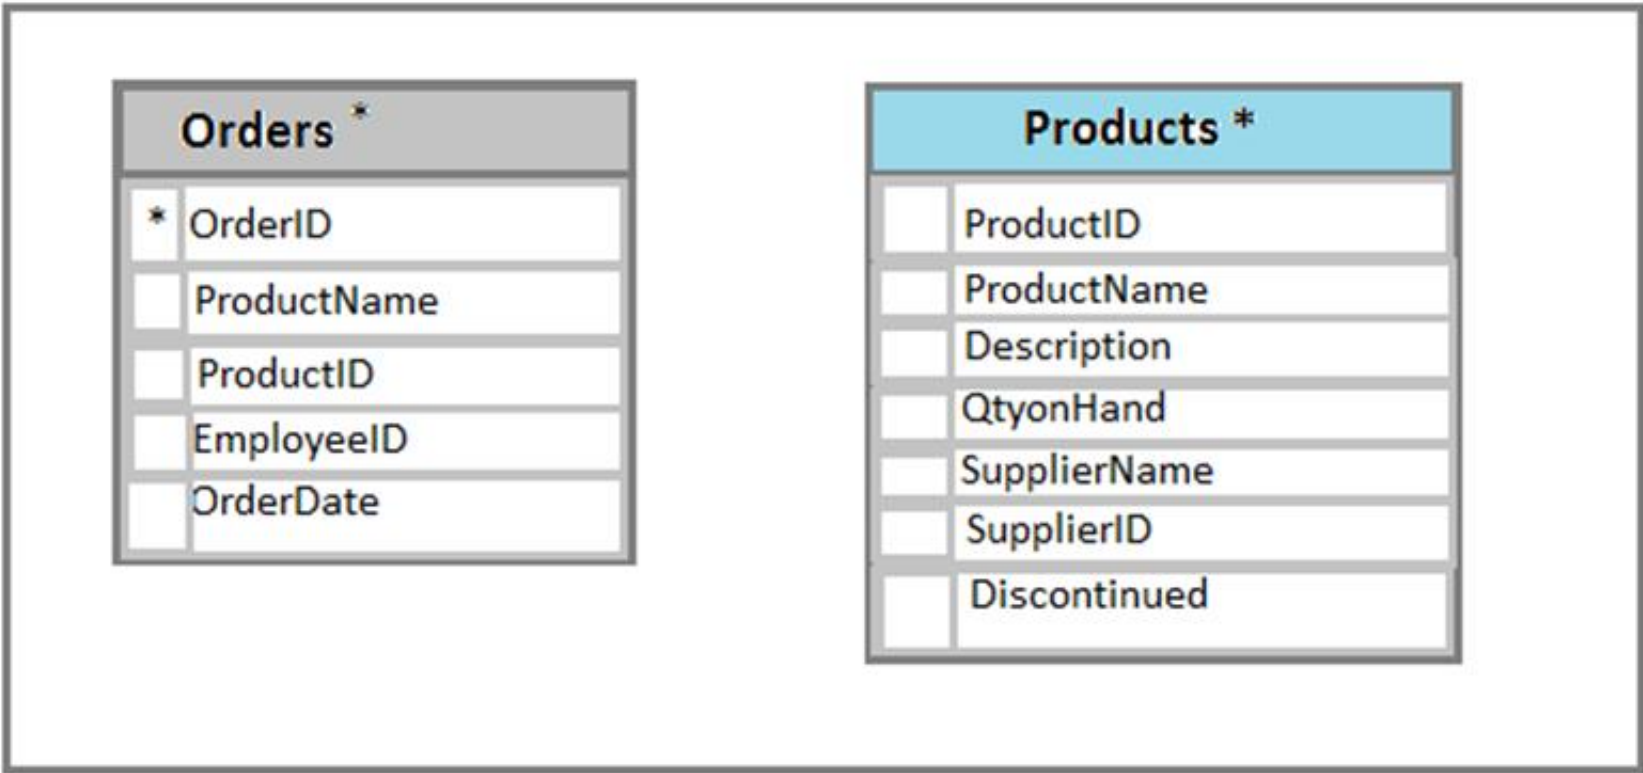
- A. Yes
- B. No

Answer: A

Explanation: References: [https://en.wikipedia.org/wiki/Check\\_constraint](https://en.wikipedia.org/wiki/Check_constraint)

NEW QUESTION 139

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 modify the database design to meet the following requirements:  
\* Rows in the Orders table must always have a valid value for the ProductID column.  
\* Rows in the Products table must not be deleted if they are part of any rows in the Orders table.  
\* All rows in both tables must be unique.  
In the table below, identify the constraint that must be configured for each table. NOTE: Make only one selection in each column.

Answer Area		Orders table	Products table
Constraint			
Check constraint on OrderID		<input type="radio"/>	<input type="radio"/>
Foreign key constraint on ProductID		<input type="radio"/>	<input type="radio"/>
Check constraint on ProductID		<input type="radio"/>	<input type="radio"/>
Foreign key constraint on OrderID		<input type="radio"/>	<input type="radio"/>

Answer:

**Explanation:** A FOREIGN KEY in one table points to a PRIMARY KEY in another table. Here the foreign key constraint is put on the ProductID in the Orders, and points to the ProductID of the Products table.

With a check constraint on the ProductID we can ensure that the Products table contains only unique rows.

References:

[http://www.w3schools.com/sql/sql\\_foreignkey.asp](http://www.w3schools.com/sql/sql_foreignkey.asp)

#### NEW QUESTION 140

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 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 are developing an application to track customer sales. You create tables to support the application. You need to create a database object that meets the following data entry requirements:

Table name	Data entry requirements
Customers	Inserts must be grouped in a batch of Transact-SQL statements.
CustomerTransactions	Inserts must be grouped in a batch of Transact-SQL statements.
Invoices	Inserts must be grouped in a batch of Transact-SQL statements. All of the data modifications made must become a permanent part of the database if data entry is successful. All data entry modifications must be erased if an error occurs.
Orders	Inserts must be grouped in a batch of Transact-SQL statements. All of the data modifications made must become a permanent part of the database if data entry is successful. All data entry modifications must be erased if an error occurs.

What should you create?

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

**Answer:** C

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

#### NEW QUESTION 145

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.Female.Customers 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  
(Segmentm 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 CustomerStatusnot equal to 1 (the active customers).

References:

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

#### NEW QUESTION 149



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 an object that meets the following requirements:

Run managed code packaged in an assembly that was created in the Microsoft.NET Framework and uploaded in Microsoft SQL Server.

Run within a transaction and roll back if a failure occurs.

Run when a table is created or modified.

What should you create?

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

**Answer:** B

**Explanation:** The common language runtime (CLR) is the heart of the Microsoft .NET Framework and provides the execution environment for all .NET Framework code. Code that runs within the CLR is referred to as managed code.

With the CLR hosted in Microsoft SQL Server (called CLR integration), you can author stored procedures, triggers, user-defined functions, user-defined types, and user-defined aggregates in managed code.

References:

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/introduction-to-sql-server-clr-integration>

## NEW QUESTION 152

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 157

You run the following Transact-SQL statement:

```
CREATE TABLE OrderLines(
    OrderLineID INT NOT NULL IDENTITY PRIMARY KEY CLUSTERED,
    OrderID INT NOT NULL,
    StockItemID INT NOT NULL,
    Description NVARCHAR(100) NOT NULL,
    Quantity INT NOT NULL,
    UnitPrice DECIMAL(18, 2) NULL
)
```

There are multiple unique OrderID values. Most of the UnitPrice values for the same OrderID are different. You need to create a single index seek query that does not use the following operators:

- \*Nested loop
- \*Sort
- \*Key lookup

- A. CREATE INDEX IX\_OrderLines\_I ON OrderLines (OrderID, UnitPrice) INCLUDE(Description, Quantity)
- B. CREATE INDEX IX\_OrderLines\_I ON OrderLines (OrderID, UnitPrice) INCLUOE(Quantity)
- C. CREATE INDEX IX\_OrderLines\_I ON OrderLines (OrderID, UnitPrice, Quantity)
- D. CREATE INDEX IX\_OrderLines\_I ON OrderLines (UnitPrice, OrderID) INCLUDE(Description, Quantity)

**Answer: A**

### NEW QUESTION 161

Case study 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.

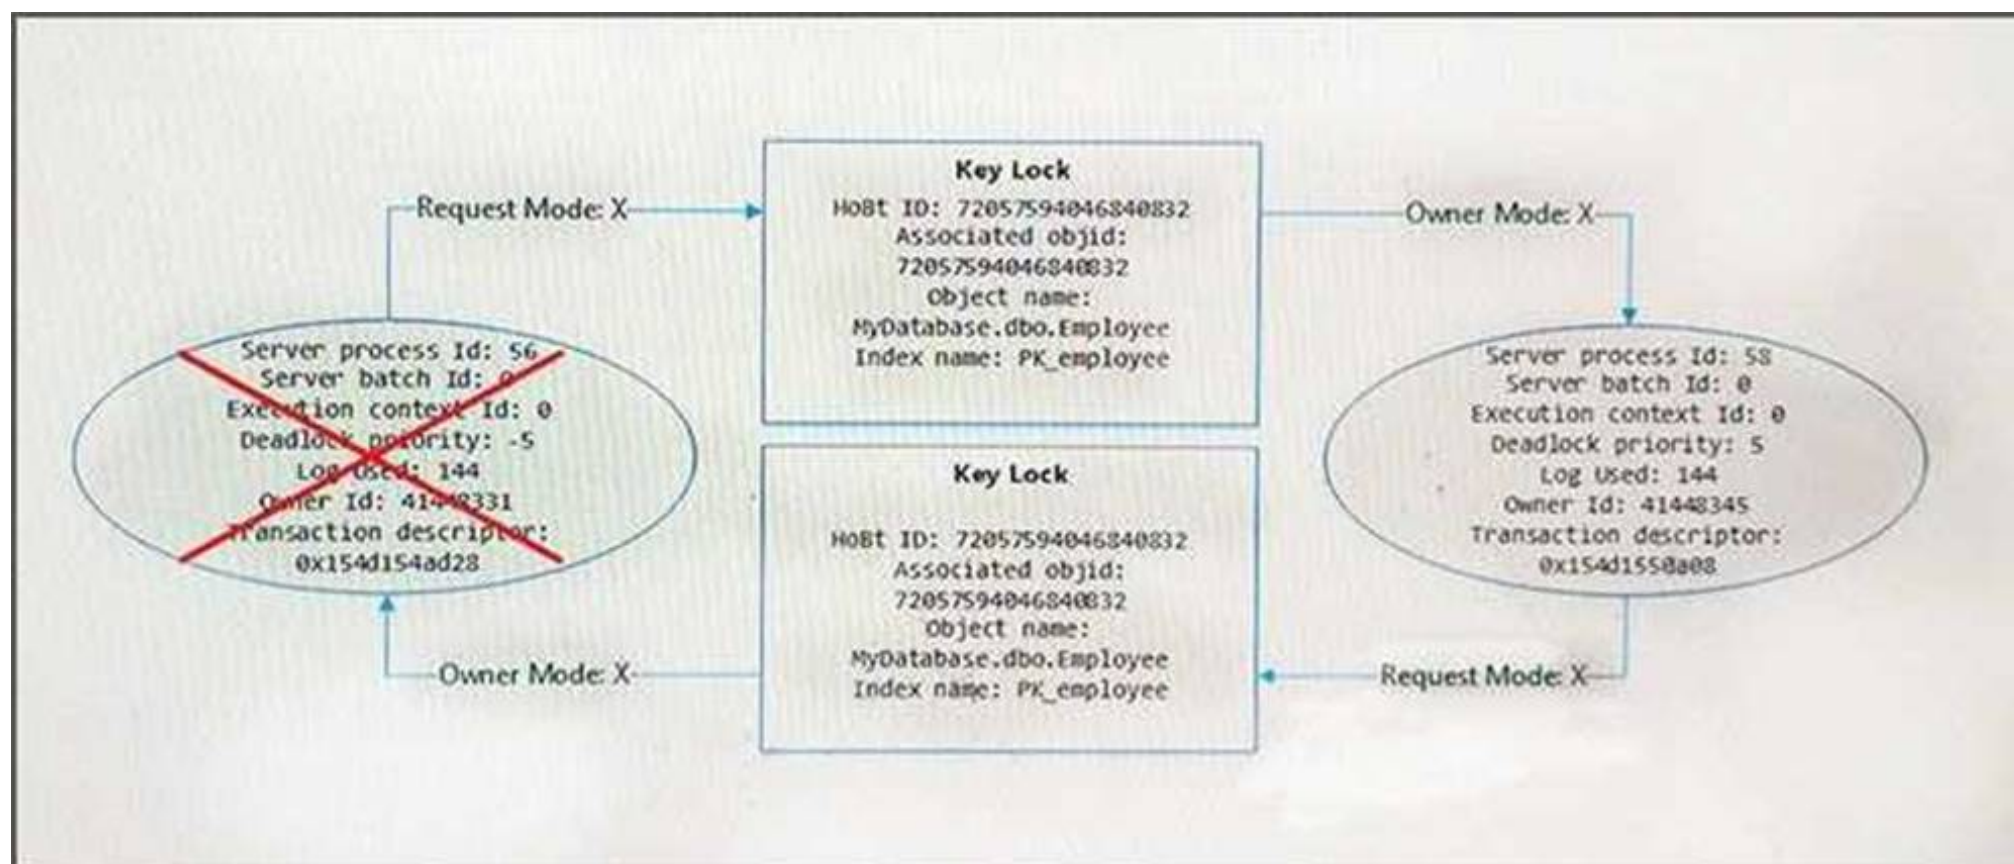
The application uses views to control access to data.

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.





You observe that the four indexes require a large amount of disk space. You must reduce the amount of disk space that the indexes are using.

You need to create a single index to replace the existing indexes.

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.

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

## Answer Area

(INCLUDE (HireDate, DepartmentId)

(LastName, FirstName)

(FirstName, LastName, DepartmentId)

INCLUDE (LastName, DepartmentId,  
EmployeeId)

CREATE INDEX IX\_New on dbo.Employee

INCLUDE (HireDate, DepartmentId  
EmployeeId)

(FirstName, LastName)

CREATE CLUSTERED INDEX IX\_New on  
dbo. Employee



Answer:

**Explanation:** Scenario: You observe that the four indexes require a large amount of disk space. You must reduce the amount of disk space that the indexes are using.

Current indexes:

```
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
```

## NEW QUESTION 164

You have a nonpartitioned table that has a single dimension. The table is named dim. Products.Projections.

The table is queried frequently by several line-of-business applications. The data is updated frequently throughout the day by two processes.

Users report that when they query data from dim.Products.Projections, the responses are slower than expected. The issue occurs when a large number of rows



are being updated.  
 You need to prevent the updates from slowing down the queries. 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.

**Answer:** B

#### NEW QUESTION 166

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 add a unique clustered index to the Employees table. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:** References: <https://msdn.microsoft.com/en-us/library/ms190457.aspx>

#### NEW QUESTION 168

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 172

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 datetime2 NOT NULL,
    DateCheclOut datetime2 NOT NULL)
```

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:** datetime2 defines a date that is combined with a time of day that is based on 24-hour clock. datetime2 can be considered as an extension of the existing datetime type that has a larger date range, a larger default fractional precision, and optional user-specified precision.  
References: <https://msdn.microsoft.com/en-us/library/bb677335.aspx>

### NEW QUESTION 177

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 @TranCount > 0
    SAVE TRANSACTION SavePoint
ELSE
    BEGIN TRANSACTION
    BEGIN TRY
        INSERT Table1 (Key1) VALUES (@Value)
        IF XACT_STATE() = 0
            COMMIT
        END TRY
        BEGIN CATCH
            IF @TranCount = 0
                ROLLBACK TRANSACTION
            ELSE IF XACT_STATE() <> -1
                ROLLBACK TRANSACTION SavePoint
        END CATCH
    END
```

### NEW QUESTION 182

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 186

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:

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:

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

When you start uspUpdateCustomerSummary, there are no active transactions. The procedure fails at line 09 due to a CHECK constraint violation on the TotalDepositAccountCount column.

What is the impact of the stored procedure on the CustomerDetails table?

- A. The value of the TotalAccountCount column decreased.
- B. The value of the TotalDepositAccountCount column is not changed.
- C. The statement that modifies TotalDepositAccountCount is excluded from the transaction.
- D. The value of the TotalAccountCount column is not changed.

**Answer:** D

### NEW QUESTION 190

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 193

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 194

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 197

You have two databases with the following settings:

Setting	Value
DELAYED_DURABILITY	ALLOWED
MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT	ON

You run the following Transact –SQL statements:

```
USE MemDb
GO
CREATE TABLE MemTable (
    Id INT IDENTITY (1, 1) PRIMARY KEY NONCLUSTERED
    DiskDbUpdateCounter INT)
WITH (MEMORY_OPTIMIZED=ON, DURABILITY=SCHEMA_ONLY)

USE DiskDb
Go
CREATE TABLE DiskTable (
    IdToUpdate INT,
    UpdateCounter INT
)
```

You need to select data from DiskTable and insert the data into MemTable. You must complete the insertion operation into MemTable as an explicit transaction without immediate durability.

Which four Transact-SQL segments should you use? 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.

### Transact-SQL statements

COMMIT TRANSACTION WITH (DELAYED\_DURABILITY = OFF)

COMMIT TRANSACTION WITH (DELAYED\_DURABILITY = ON)

IF OBJECT\_ID ('tempdb..#DiskTable') IS NOT NULL  
DROP TABLE #DiskTable  
SELECT \* INTO #DiskTable from DiskDb.DiskTable

BEGIN TRANSACTION

UPDATE T SET DiskDbUpdateCounter =  
DiskDbUpdateCounter + 1  
FROM MemDb.MemTable T, DiskDb.DiskTable S  
WHERE S.IdToUpdate = T.Id

UPDATE T SET DiskUpdateCounter =  
DiskDbUpdateCounter + 1  
FROM MemDb.MemTable T, #DiskTable S  
WHERE S.IdToUpdate = T.Id

ALTER DATABASE DiskDb SET DELAYED\_DURABILITY = FORCED

### Answer Area



Answer:

**Explanation:** Box 1: BEGIN TRANSACTION

Box 2: UPDATE ... #DiskTable

Box 3: IF... SELECT INTO ...#DiskTable Box 4: .. DELAYED\_DURABILITY = ON

The COMMIT syntax is extended so you can force delayed transaction durability. If DELAYED\_DURABILITY is DISABLED or FORCED at the database level (see above) this COMMIT option is ignored.

Syntax:

COMMIT [ { TRAN | TRANSACTION } ] [ transaction\_name | @tran\_name\_variable ] [ WITH ( DELAYED\_DURABILITY = { OFF | ON } ) ]

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/logs/control-transaction-durability?view=sql-server-20>

### NEW QUESTION 199

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 200

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 205

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.

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

### Answer Area

CREATE **FUNCTION** ▼ Sales.YIDSalesByPerson

(@SalesPersonID int, @minYIDSales money)

**RETURNS TABLE** ▼

AS

RETURN (SELECT TOP(@SalesPersonID) BusinessEntityID, SalesYID

FROM Sales.SalesPerson

WHERE SalesYID > @minYIDSales

ORDER BY SalesYID desc);

**Answer:**

#### Explanation:

From question: 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.

CREATE VIEW (Transact-SQL) creates a virtual table whose contents (columns and rows) are defined by a query. Use this statement to create a view of the data in one or more tables in the database.

SCHEMABINDING binds the view to the schema of the underlying table or tables. When SCHEMABINDING is specified, the base table or tables cannot be modified in a way that would affect the view definition.

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

### NEW QUESTION 209

.....

## 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