

70-762 Dumps

Developing SQL Databases (beta)

<https://www.certleader.com/70-762-dumps.html>



NEW QUESTION 1

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

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

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

Answer: D

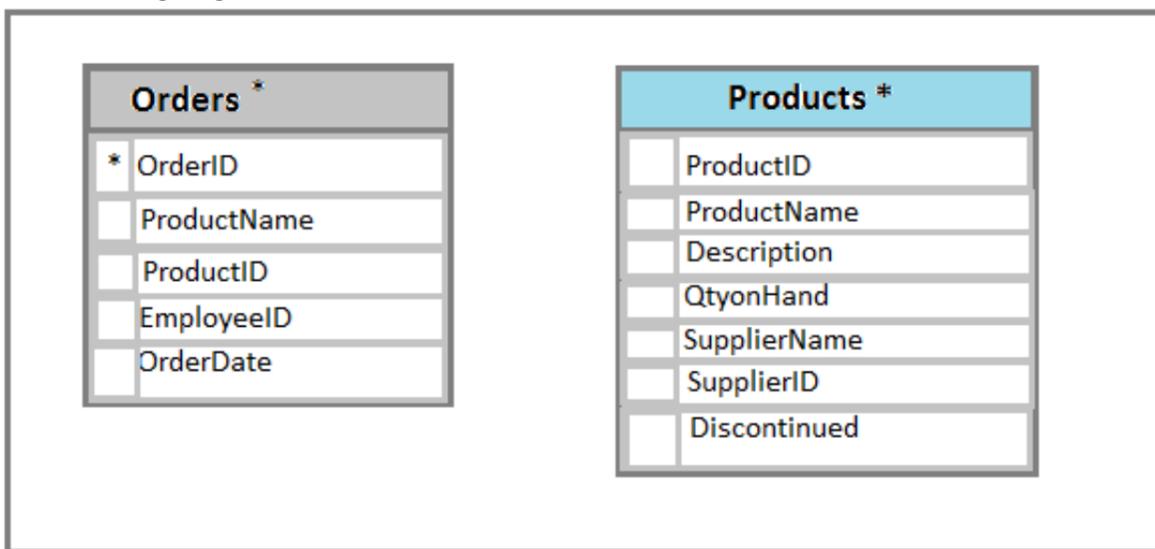
Explanation:

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

NEW QUESTION 2

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

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



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

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

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	Transact-SQL segment
Handle errors	<div style="background-color: #cccccc; height: 15px; margin-bottom: 5px;"></div> Try/Parse Select @@error <hr/> Begin Tran/Rollback Tran <hr/> Try/Catch*
Display error message	<div style="background-color: #cccccc; height: 15px; margin-bottom: 5px;"></div> ERROR MESSAGE() <hr/> PRINT <hr/> RAISERROR <hr/> RETURN

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Using TRY...CATCH in Transact-SQL

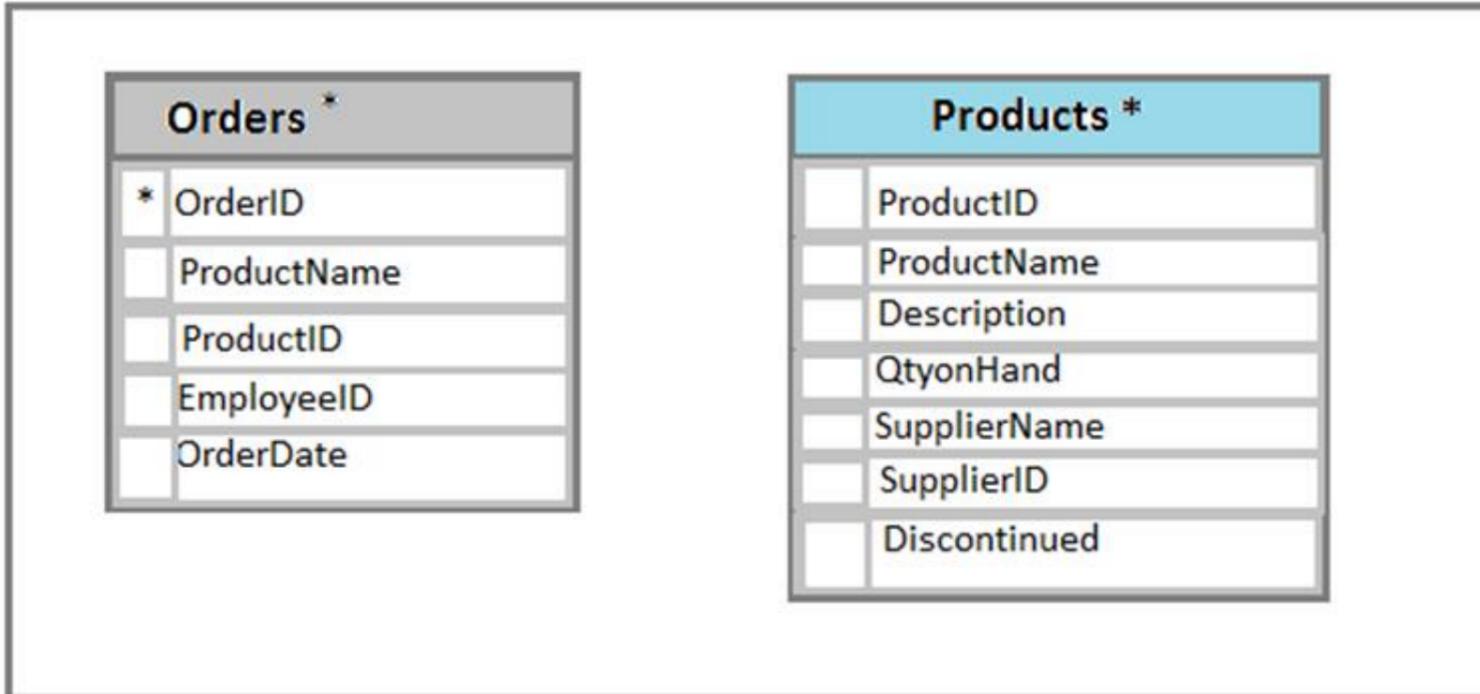
Errors in Transact-SQL code can be processed by using a TRY...CATCH construct. TRY...CATCH can use the following error function to capture error information: ERROR_MESSAGE() returns the complete text of the error message. The text includes the values supplied for any substitutable parameters such as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

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"/>

- A. Mastered
- B. Not Mastered

Answer: A

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 4

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.)

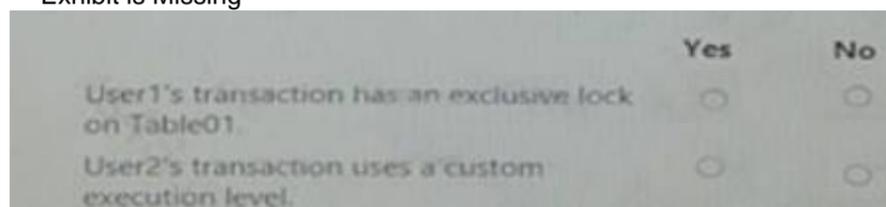
Exhibit is Missing

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.)

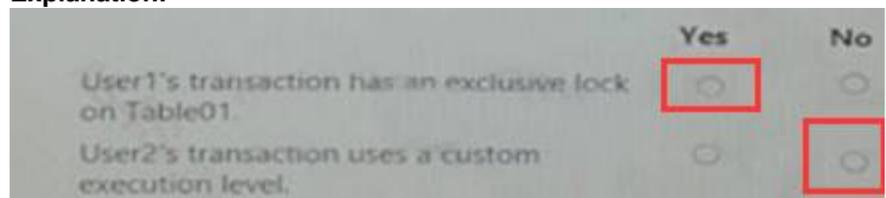
Exhibit is Missing



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 5

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

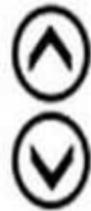
NEW QUESTION 6

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



- A. Mastered
- B. Not Mastered

Answer: A

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 7

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 written a transaction and roll back if a future 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 8

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 9

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 na 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 10

You have a database named Sales.

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

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

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

Answer Area

CREATE TABLE Customer

(

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

▼

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

MASKED WITH (FUNCTION = 'default()')

MASKED WITH (FUNCTION = 'random(1,100)')

MobileNO varchar(10)

▼

MASKED WITH (FUNCTION = 'partial(3, "XXXXXXX", 0)')

MASKED WITH (FUNCTION = 'default()')

MASKED WITH (FUNCTION = 'random(1,3)')

NULL,

Email varchar(100)

▼

MASKED WITH (FUNCTION = 'email()')

MASKED WITH (FUNCTION = 'default()')

MASKED WITH (FUNCTION = 'partial(1, "XXXXXXX", 4)')

MASKED WITH (FUNCTION = 'partial(1, "@", 4)')

NULL,

);

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

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

PhoneNumber varchar(10) MASKED WITH (FUNCTION = 'partial(5,"XXXXXX",0)') Box 3: MASKED WITH (FUNCTION = 'email()')
The Email masking method which exposes the first letter of an email address and the constant suffix ".com", in the form of an email address. .aXXX@XXX.com.
Example definition syntax: Email varchar(100) MASKEDWITH (FUNCTION = 'email()') NULL References: <https://msdn.microsoft.com/en-us/library/mt130841.aspx>

NEW QUESTION 10

Note: this question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question 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 11

You use Query Store to optimize a query in a database. The query has two execution plans:

Plan 2 is shown in the Plan 2 Execution Plan exhibit.

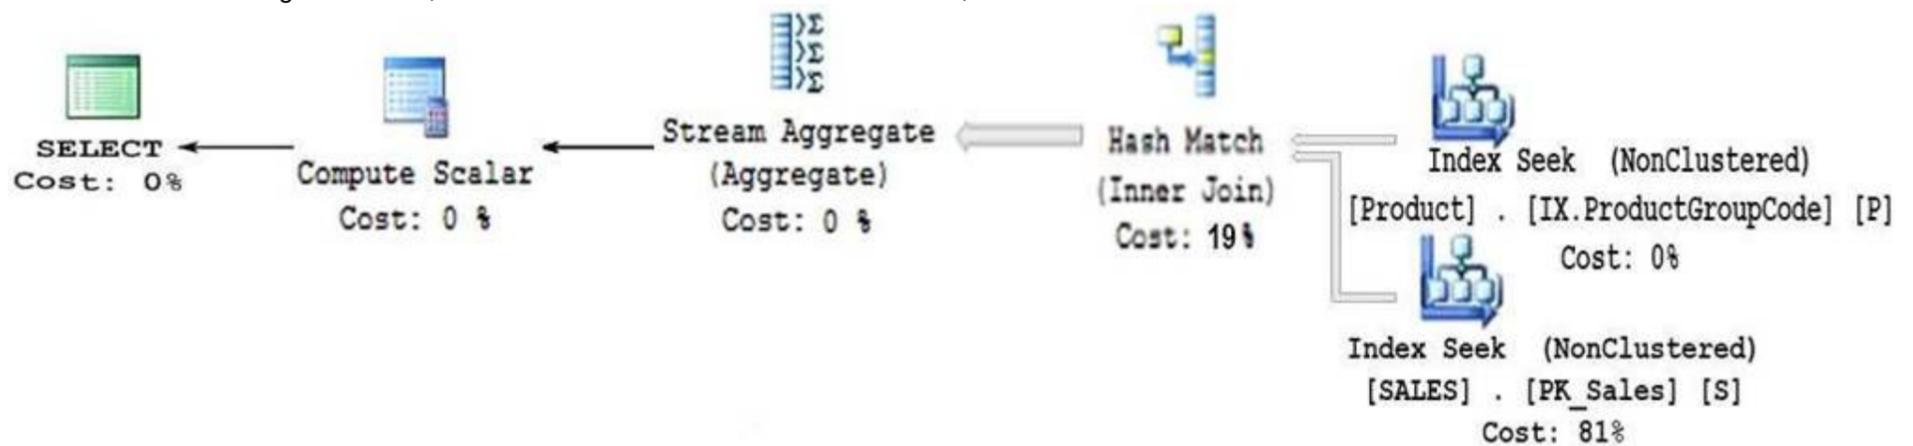
Plan 10 is shown in the Plan 10 Execution Plan exhibit.

You create an index at 22:24 based on the missing index suggestion in Plan 2.

The average duration statistics for the query is shown in the Tracked Queries exhibit. (Click the Exhibit button.)

You need to analyze the operators in the two execution plans.

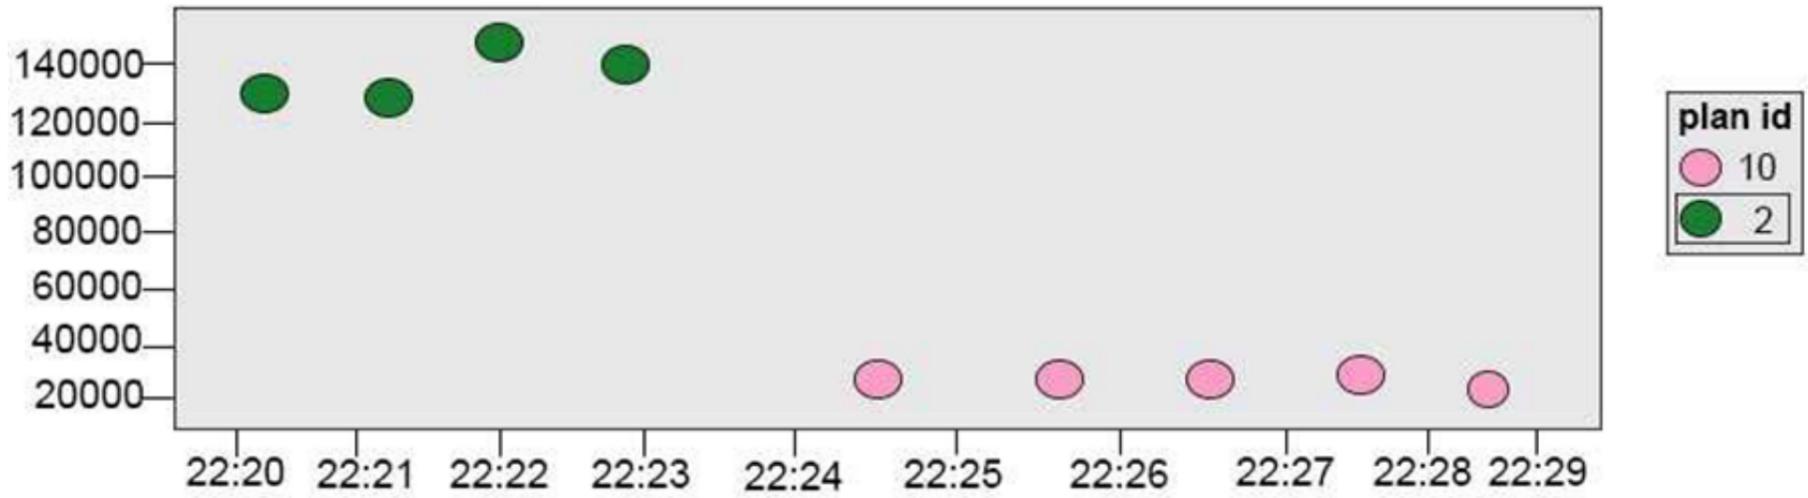
For each of the following statements, select Yes if the statement is true. Otherwise, select No. Plan 2 Execution Plan



Plan 10 Execution Plan



Tracked Queries



Answer Area

	Yes	No
The Hash Match operator in Plan 10 processed more data than the operator in Plan 2	<input type="radio"/>	<input type="radio"/>
The Clustered Index Scan operator in Plan 2 was less efficient than the Nonclustered Index Seek operator in Plan 10	<input type="radio"/>	<input type="radio"/>
Query throughput will increase by using Plan 10 instead of Plan 2	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

	Yes	No
The Hash Match operator in Plan 10 processed more data than the operator in Plan 2	<input checked="" type="radio"/>	<input type="radio"/>
The Clustered Index Scan operator in Plan 2 was less efficient than the Nonclustered Index Seek operator in Plan 10	<input checked="" type="radio"/>	<input type="radio"/>
Query throughput will increase by using Plan 10 instead of Plan 2	<input type="radio"/>	<input checked="" type="radio"/>

NEW QUESTION 13

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 16

You have a database named MyDatabase. You must monitor all the execution plans in XML format by using Microsoft SQL Trace. The trace must meet the following requirements:

- Capture execution plans only for queries that run the MyDatabase database.
- Filter out plans with event duration of less than or equal to 100 microseconds.
- Save trace results to a disk on the server. You need to create the trace.

In which order should you arrange the Transact-SQL segments to develop the solution? To answer, move all 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.

DECLARE @traceEventId int = 122;

DECLARE @traceColumnIdForTextData int = 1; DECLARE @durationFilter bigint = 100 DECLARE @databaseID int;

SELECT @databaseId = DB_ID('MyDatabase');

Transact-SQL segments

```
EXEC sp_trace_setfilter @TraceID, 13, 0, 2, @duration-
Filter;
EXEC sp_trace_setfilter @TraceID, 3, 0, 0, @databaseId;
```

```
EXEC sp_trace_setevent @TraceID, @traceEventId, 1, 1;
EXEC sp_trace_setevent @TraceID, @traceEventId, 12, 1;
EXEC sp_trace_setevent @TraceID, @traceEventId, 13, 1;
```

```
EXEC sp_trace_setstatus @TraceID, 1;
```

```
EXEC sp_trace_create @TraceID OUTPUT, 2, @fileName, @max-
filesize, NULL @fileCount;
```

```
DECLARE @ traceID int;
DECLARE @maxFileSize bigint = 20;
DECLARE @fileCount INT = 15;
DECLARE @fileName NVARCHAR(245) = 'D:\SQL Trace\TraceRe-
sult;
DECLARE @traceEventId int = 122;
DECLARE @traceColumnIdForTextData int = 1;
DECLARE @durationFilter bigint = 100;
DECLARE @databaseID int;
('SELECT @databaseId = DB_ID('MyDatabase');
```

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The following system stored procedures are used to define and manage traces:

* sp_trace_create is used to define a trace and specify an output file location as well as other options that I'll cover in the coming pages. This stored procedure returns a handle to the created trace, in the form of an integer trace ID.

* sp_trace_setevent is used to add event/column combinations to traces based on the trace ID, as well as to remove them, if necessary, from traces in which they have already been defined.

* sp_trace_setfilter is used to define event filters based on trace columns.

* sp_trace_setstatus is called to turn on a trace, to stop a trace, and to delete a trace definition once you're done with it. Traces can be started and stopped multiple times over their lifespan.

References: <https://msdn.microsoft.com/en-us/library/cc293613.aspx>

NEW QUESTION 17

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 22

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 Enterprise Edition interface. At the top, there's a toolbar with options like 'New Query'. Below that, a tab shows 'genemi4f\sql2016rc...xec_ES_2: Live Data'. A table displays 'Displaying 30 Events' with columns for 'name' and 'timestamp'. Below the table, an event 'query_post_execution_showplan' is selected, and the 'Query Plan' tab is active. The query plan shows a 'Nested Loops (Left Outer Join)' node highlighted with a red box. A red arrow points from this node to a detailed properties pane on the right. The properties pane is titled 'Clustered Index Seek (Clustered)' and lists various metrics such as 'Physical Operation', 'Logical Operation', 'Actual Execution Mode', 'Estimated Execution Mode', 'Storage', 'Actual Number of Rows', 'Actual Number of Batches', 'Estimated Operator Cost', 'Estimated I/O Cost', 'Estimated Subtree Cost', 'Estimated CPU Cost', 'Estimated Number of Executions', 'Number of Executions', 'Estimated Number of Rows', 'Estimated Row Size', 'Actual Rebinds', 'Actual Rewinds', 'Ordered', and 'Node ID'. It also shows the 'Object' as '[InMemTest2].[sys].[sysclsobjs].[clst] [s]' and the 'Output List' as '[InMemTest2].[sys].[sysclsobjs].name'. The 'Seek Predicates' section includes 'Seek Keys[1]: Prefix: [InMemTest2].[sys].[sysclsobjs].class, [InMemTest2].[sys].[sysclsobjs].id = Scalar Operator((50)), Scalar Operator([InMemTest2].[sys].[sysclsobjs].[nsid] as [s], [nsid])'.

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

NEW QUESTION 26

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,
    DateCheckedIn datetimeoffset(0) NOT NULL,
    DateCheckedOut 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.

Syntax: datetimeoffset [(fractional seconds precision)]

For the 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 31

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.

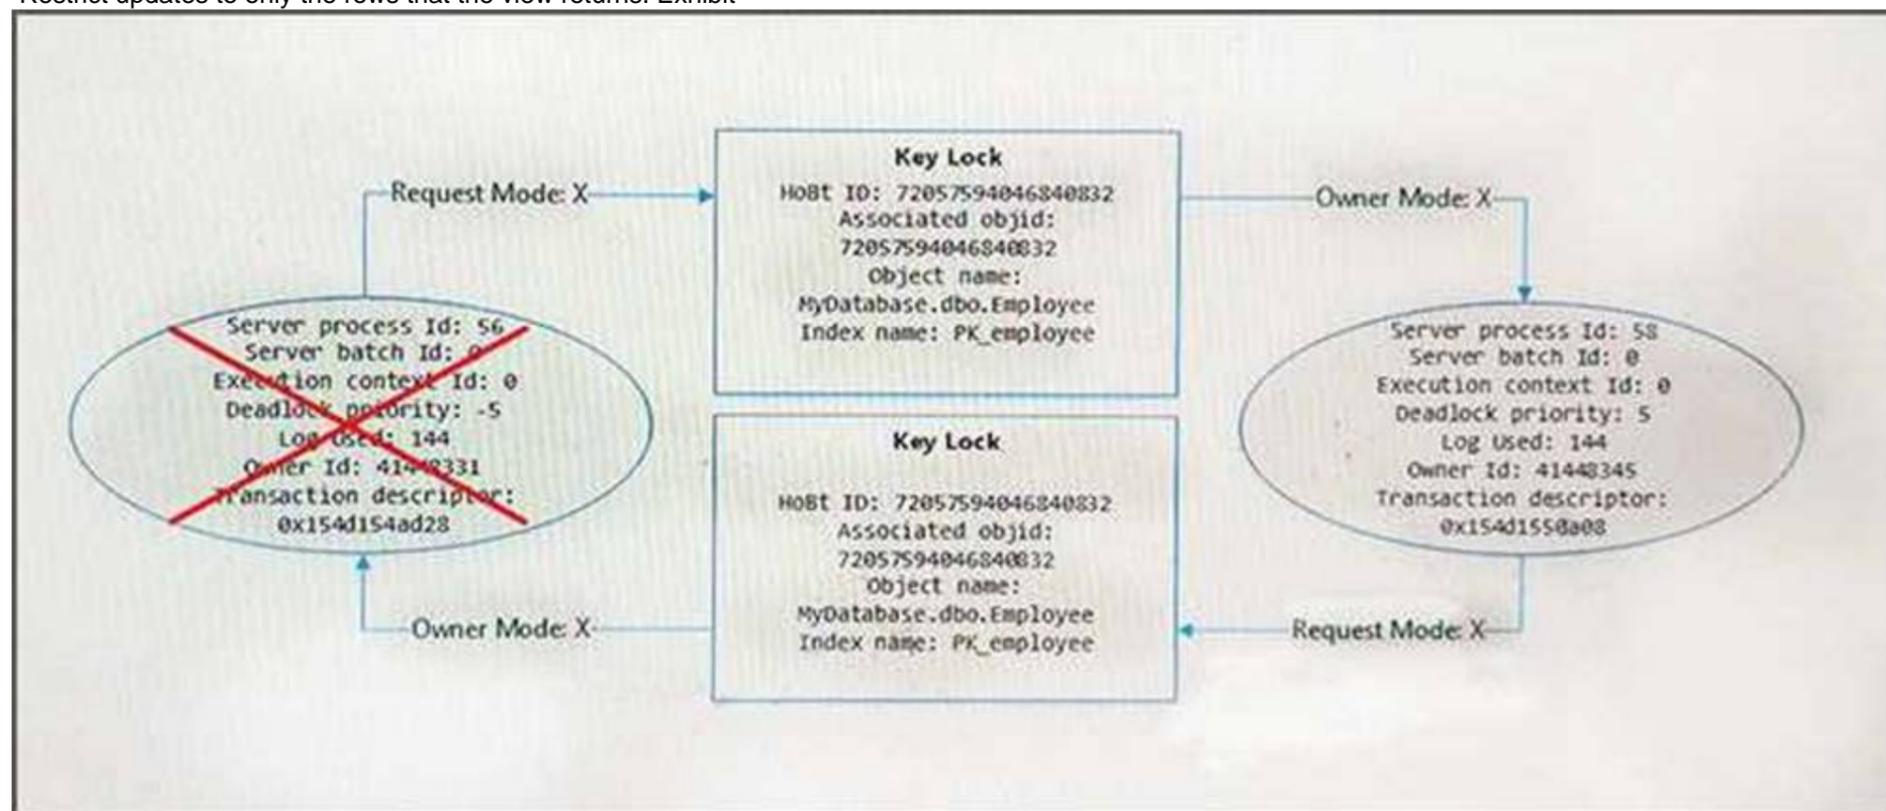
UspB: This stored procedure updates only the EmployeePayRate column.

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 36

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

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 index on TBL1.
- F. Create a nonclustered columnstore index on TBL1. Create nonclustered index on TBL2.
- G. Drop existing indexes on TBL1 and then create a clustered columnstore index on TBL1.
- H. Create a nonclustered columnstore index on TBL1. Make no changes to TBL2.

Answer: B

NEW QUESTION 39

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

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

Transact-SQL segments

```
RETURNS TABLE AS
```

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

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

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

```
RETURNS [Money] AS
```

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



Answer area

```
CREATE FUNCTION dbo.GetPrice (@ProductId int)
```

Transact-SQL segment

Transact-SQL segment

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References: [https://technet.microsoft.com/en-us/library/ms189294\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms189294(v=sql.105).aspx)

NEW QUESTION 44

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

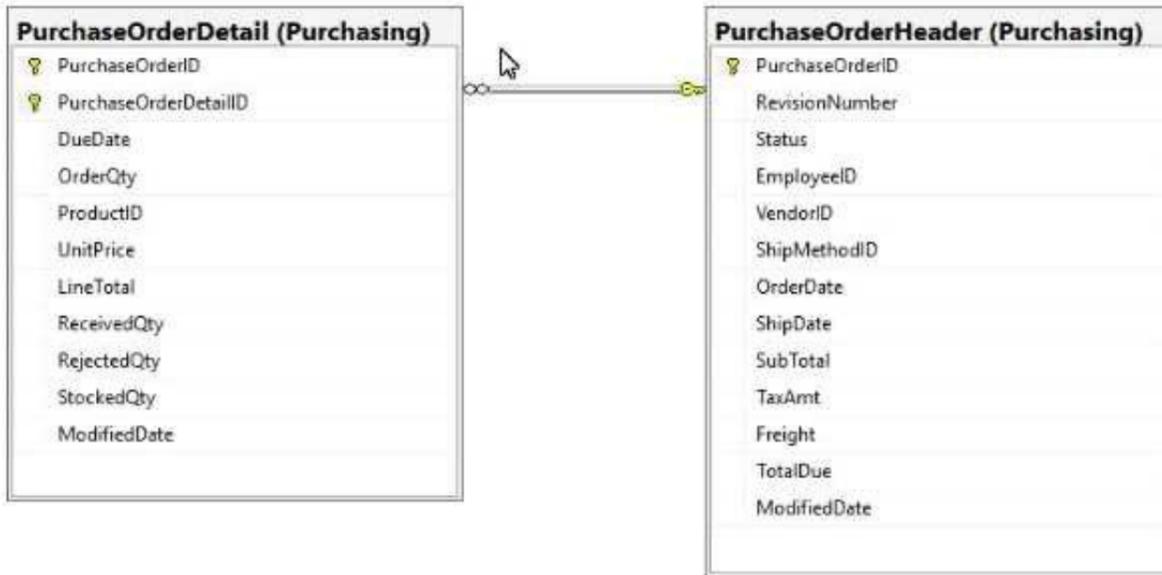
Answer: C

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 47

You manage a database that includes the tables shown in the exhibit (Click the Exhibit button.)



You plan to create a DML trigger that reads the value of the LineTotal column for each row in the PurchaseOrderDetail table. The trigger must add the value obtained to the value in the SubTotal column of the PurchaseOrderHeader table. You need to organize the list to form the appropriate Transact-SQL statement. Which five 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.

Transact-SQL segments

- CREATE TRIGGER NewPODetail
- UPDATE PurchaseOrderDetail
- UPDATE PurchaseOrderHeader
SET SubTotal = SubTotal + LineTotal
FROM inserted
- UPDATE PurchaseOrderDetail
SET SubTotal = SubTotal + LineTotal
FROM inserted
- ON PurchaseOrderDetail
- UPDATE PurchaseOrderHeader
SET SubTotal = LineTotal
FROM inserted
- AFTER INSERT AS
- WHERE PurchaseOrderHeader.PurchaseOrderID =
inserted.PurchaseOrderID

Answer area

⬆
⬇

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer area

- 1 AFTER INSERT AS
- 2 CREATE TRIGGER NewPODetail
- 3 ON PurchaseOrderDetail

NEW QUESTION 52

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 need to create a stored procedure that updates the Customer, CustomerInfo, OrderHeader, and OrderDetails tables in order.

You need to ensure that the stored procedure:

- Runs within a single transaction.
- Commits updates to the Customer and CustomerInfo tables regardless of the status of updates to the OrderHeader and OrderDetail tables.
- Commits changes to all four tables when updates to all four tables are successful. Solution: You create a stored procedure that includes the following Transact-SQL segment:

```
BEGIN TRY
    BEGIN TRAN
        UPDATE Customer ...
        UPDATE CustomerInfo ...
        UPDATE OrderHeader ...
        UPDATE OrderDetail ...
    COMMIT TRAN
END TRAN
BEGIN CATCH
    IF XACT_STATE() = 1
        ROLLBACK TRAN
END CATCH
```

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

References:

<http://stackoverflow.com/questions/11444923/stored-procedure-to-update-multiple-tables>

NEW QUESTION 54

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 57

You have a view that includes an aggregate.

You must be able to change the values of columns in the view. The changes must be reflected in the tables that the view uses.

You need to ensure that you can update the view. What should you create?

- A. table-valued function
- B. a schema-bound view
- C. a partitioned view
- D. a DML trigger

Answer: B

Explanation:

When you use the SchemaBinding keyword while creating a view or function you bind the structure of any underlying tables or views. It means that as long as that schemabound object exists as a schemabound object (ie you don't remove schemabinding) you are limited in changes that can be made to the tables or views that it refers to.

References: <https://sqlstudies.com/2014/08/06/schemabinding-what-why/>

NEW QUESTION 60

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 ProductCode = 'CD'
```

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

```
CREATE CLUSTERED INDEX PK_Account ON Account(ProductCode); Does the solution meet the goal?
```

- A. Yes
- B. No

Answer: B

Explanation:

We need an index on the productCode column as well.

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

NEW QUESTION 65

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 create a stored procedure named spDeleteCategory to delete records in the database. The stored procedure must meet the following requirements:

- * Delete records in both the BlogEntry and BlogCategory tables where CategoryId equals parameter @CategoryId.

* Avoid locking the entire table when deleting records from the BlogCategory table.

* If an error occurs during a delete operation on either table, all changes must be rolled back, otherwise all changes should be committed.

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

Answer Area

```

CREATE PROCEDURE spDeleteCategory
(@CategoryID int)
AS
BEGIN
    SET NOCOUNT ON:
    [Box 1]
    SET IMPLICIT_TRANSACTIONS ON
    SET IMPLICIT_TRANSACTIONS OFF
    SET TRANSACTION ISOLATION LEVEL READ COMMITTED
    SET TRANSACTION ISOLATION LEVEL SNAPSHOT
    BEGIN TRY
        DELETE FROM BlogEntry WHERE CategoryID = @CategoryID;
        ...
        DELETE FROM BlogCategory
        WITH ( [Box 2] ) WHERE CategoryId = @CategoryID;
        IF @@TRANCOUNT > 0 [Box 3] TRANSACTION;
        BE
    END TRY
    BEGIN CATCH
        IF @@TRANCOUNT > 0 [Box 4] TRANSACTION;
    END
END
    
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: SET TRANSACTION ISOLATION LEVEL READ COMMITTED

You can minimize locking contention while protecting transactions from dirty reads of uncommitted data modifications by using either of the following:

- * The READ COMMITTED isolation level with the READ_COMMITTED_SNAPSHOT database option set ON.
- * The SNAPSHOT isolation level.

With ROWLOCK we should use READ COMMITTED Box 2: ROWLOCK

Requirement: Avoid locking the entire table when deleting records from the BlogCategory table

ROWLOCK specifies that row locks are taken when page or table locks are ordinarily taken. When specified in transactions operating at the SNAPSHOT isolation level, row locks are not taken unless ROWLOCK is combined with other table hints that require locks, such as UPDLOCK and HOLDLOCK.

Box 3: COMMIT

Box 4: ROLLBACK

NEW QUESTION 66

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

▼

- MyInputTable#
- MyInputTable@
- MyDatatype@

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 67

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 71

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

- A. Mastered
- B. Not Mastered

Answer: A

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 73

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

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

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

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

You frequently run the following Transact-SQL select statements:

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 78

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 83

You are developing an application that connects to a database. The application runs the following jobs:

Job	Transact-SQL statement	Description
JobA	Exec uspDeletePrevRecords	The stored procedure deletes all records from a table named tblBalanceTransactions that were created before the current month by using a single DELETE statement. Approximately 10 million records are deleted each time you run this stored procedure
JobB	Exec uspUpdateCurRecords	This stored procedure updates records in the tblBalanceTransaction table that were created in the current month. Only a few hundred records are updated each time you run this stored procedure.

The READ_COMMITTED_SNAPSHOT database option is set to OFF, and auto-content is set to ON. Within the stored procedures, no explicit transactions are defined.

If JobB starts before JobA, it can finish in seconds. If JobA starts first, JobB takes a long time to complete. You need to use Microsoft SQL Server Profiler to determine whether the blocking that you observe in JobB is caused by locks acquired by JobA. Which trace event class in the Locks event category should you use?

- A. LockAcquired
- B. LockCancel
- C. LockDeadlock
- D. LockEscalation

Answer: A

Explanation:

The Lock:Acquired event class indicates that acquisition of a lock on a resource, such as a data page, has been achieved.

The Lock:Acquired and Lock:Released event classes can be used to monitor when objects are being locked, the type of locks taken, and for how long the locks were retained. Locks retained for long periods of time may cause contention issues and should be investigated.

NEW QUESTION 84

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.dm_resource_usage
 - E. sys.dm_db_resource_stats

Answer: BC

NEW QUESTION 85

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 87

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

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

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

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

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

There are no foreign key relationships between TBL1 and TBL2.

You need to minimize the amount of time required for the two queries to return records from the tables. What should you do?

- A. Create clustered indexes on TBL1 and TBL2.
- B. Create a clustered index on TBL1 Create a nonclustered index on tbl2 and add the most frequently queried columns as included columns.
- C. Create a nonclustered index on tbl2 only.
- D. Create unique constraints on both TBL1 and TBL2. Create a partitioned view that combines columns from TBL1 and TBL2.
- E. Drop existing indexes on TBL1 and then create a clustered columnstore index
- F. Create a nonclustered columnstore index on TBL1. Create a nonclustered index on TBL2.
- G. Drop existing indexes on TBL1 and then create a clustered columnstore index
- 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 90

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
```

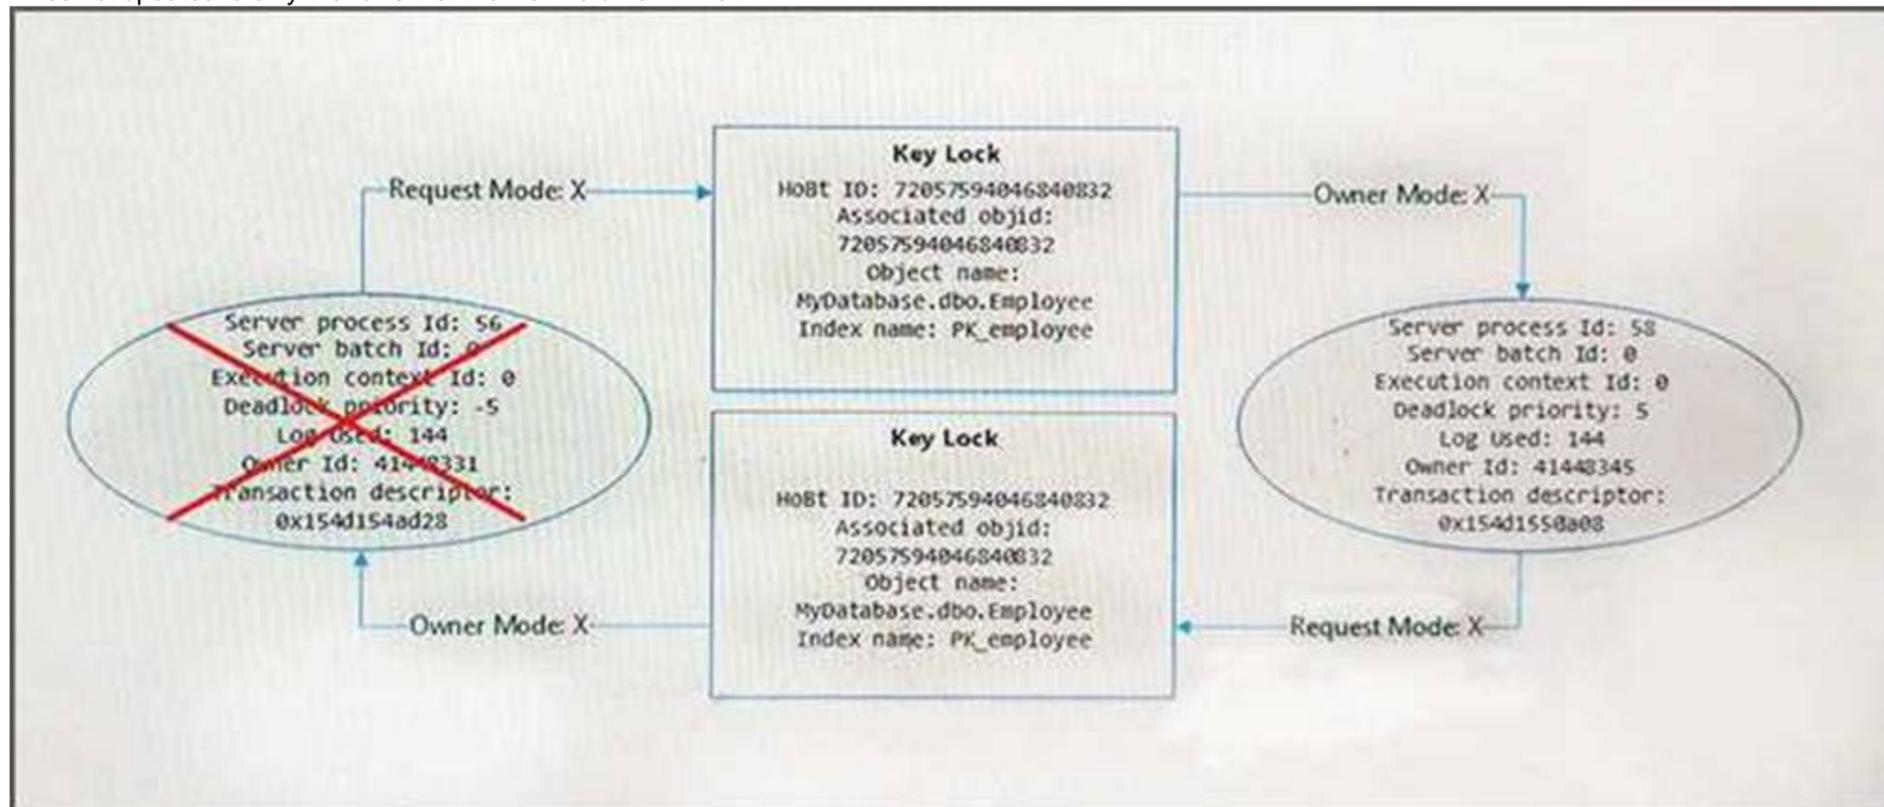
Application

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"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 92

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

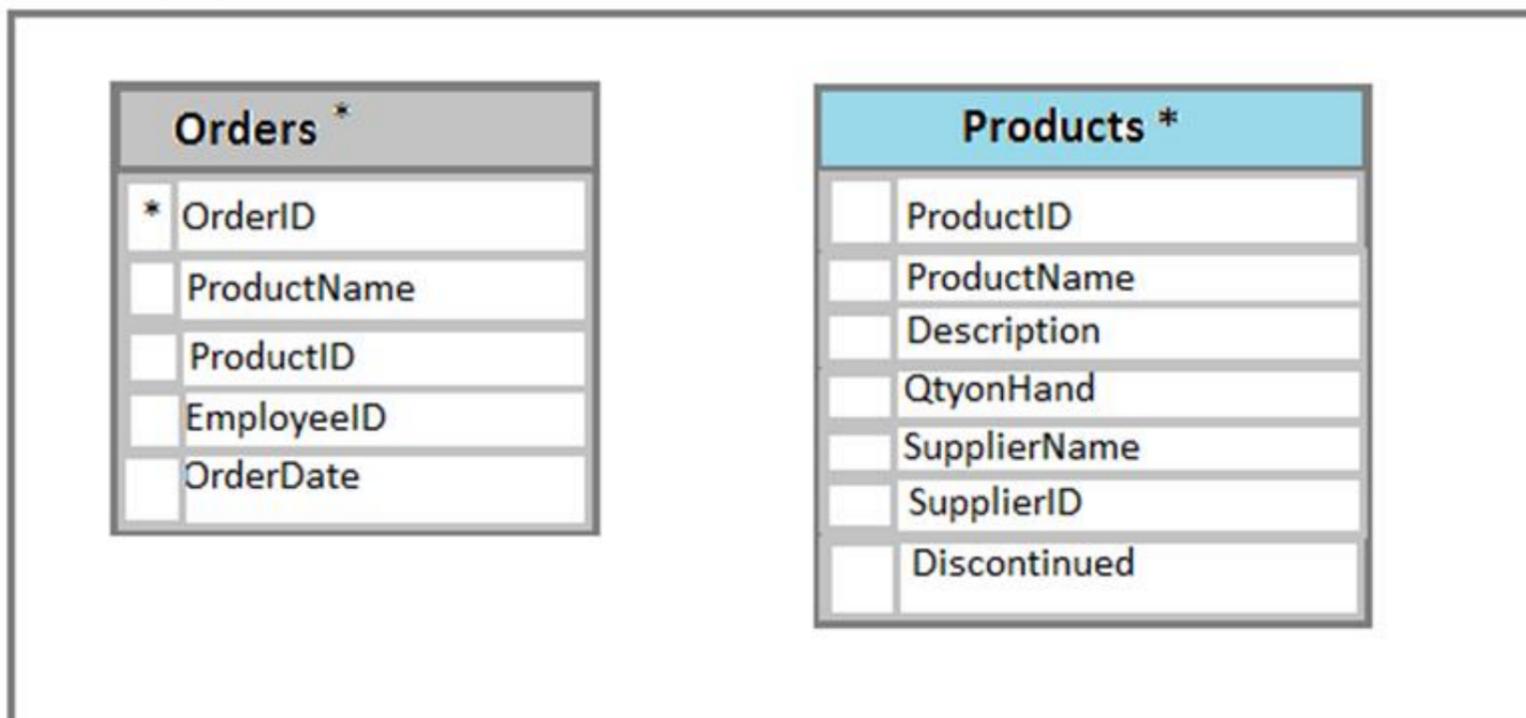
Explanation:

References: https://en.wikipedia.org/wiki/Check_constraint

NEW QUESTION 93

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

Constraint	Orders table	Products table
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"/>

- A. Mastered
- B. Not Mastered

Answer: A

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

NEW QUESTION 97

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. 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.
Invoices	All data entry modifications must be erased if an error occurs. 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.
Orders	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 99

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
(Segment, ContryCode, PhoneNumber)

WHERE CustomerID = 1;

CREATE NONCLUSTERED INDEX IX_Customer_CreatedDate

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

CREATE CLUSTERED INDEX IX_Customer_CreatedDate
```



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Clustered Index

With the same size of keys, the nonclustered indexes need more space than clustered indexes. Box 2, Box 3:

Include the CustomerStatus column in the index, and only when CustomerStatus not equal to 1 (the active customers).

References:

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

NEW QUESTION 101

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 to following requirements:

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

Run written 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 103

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

- A. Mastered
- B. Not Mastered

Answer: A

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 108

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

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

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

Solution: You create a stored procedure in the sales schema that selects the data from the sales. orders table.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 109

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 110

You are analyzing the performance of a database environment.

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

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

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

Transact-SQL segments

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

Answer Area

```

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

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

SELECT      TOP (5000) Transact-SQL statement

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

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: sys.db_db_missing_index_group_stats

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

Box 2: group_handle

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

SELECT migs.group_handle, mid.*

FROM sys.dm_db_missing_index_group_stats AS migs INNER JOIN sys.dm_db_missing_index_groups AS mig ON (migs.group_handle = mig.index_group_handle) INNER JOIN sys.dm_db_missing_index_details AS mid ON (mig.index_handle = mid.index_handle)

WHERE migs.group_handle = 24;

Box 3: sys.db_db_missing_index_group_stats

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

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

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

SELECT TOP 10 *

FROM sys.dm_db_missing_index_group_stats

ORDER BY avg_total_user_cost * avg_user_impact * (user_seeks + user_scans)DESC;

NEW QUESTION 111

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 114

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 118

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>

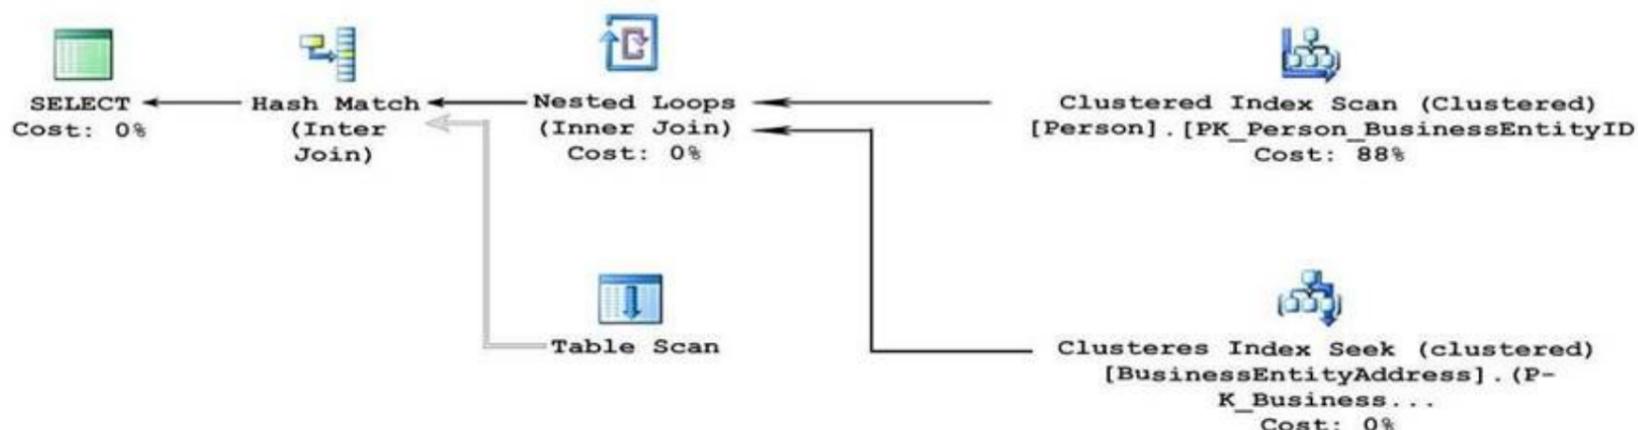
NEW QUESTION 121

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

You have a database named DB1. There is no memory-optimized filegroup in the database. You run the following query:

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

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



Users frequently run the same query with different values for the local variable @lastName. The table named Person is persisted on disk.

You need to create an index on the Person.Person table that meets the following requirements:

- All users must be able to benefit from the index.
- FirstName must be added to the index as an included column. What should you do?

- A. Create a clustered index on the table.
- B. Create a nonclustered index on the table.

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

Answer: B

Explanation:

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

NEW QUESTION 124

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution 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 change the transaction log file size to expand dynamically in increments of 200 MB. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 127

You are designing a solution for a company that operates retail stores. Each store has a database that tracks sales transactions. You create a summary table in the database at the corporate office. You plan to use the table to record the quantity of each product sold at each store on each day. Managers will use this data to identify reorder levels for products.

Every evening, stores must transmit sales data to the corporate office. The data must be inserted into the summary table that includes the StoreID, ProductID, QtySold, TotalSales, and DateSold columns.

You need to prevent duplicate rows in the summary table. Each row must uniquely identify the store that sold the product and the total amount sold for that store on a specific date.

What should you include in your solution?

- A. Create a unique constraint.
- B. Create a foreign key constraint to the storeID column in each of the store tables.
- C. Create a rule and bind it to the storeID column.
- D. Create a check constraint.

Answer: B

NEW QUESTION 128

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 need to create a stored procedure that updates the Customer, CustomerInfo, OrderHeader, and OrderDetails tables in order.

You need to ensure that the stored procedure:

- Runs within a single transaction.
- Commits updates to the Customer and CustomerInfo tables regardless of the status of updates to the OrderHeader and OrderDetail tables.
- Commits changes to all four tables when updates to all four tables are successful. Solution: You create a stored procedure that includes the following Transact-SQL segment:

```
DECLARE @CustomerComplete bit = 0
BEGIN TRY
    BEGIN TRAN
        UPDATE Customer ...
        UPDATE CustomerInfo ...

        SET @CustomerComplete = 1

        SAVE TRAN TR1

        UPDATE OrderHeader ...
        UPDATE OrderDetail ...
    COMMIT TRAN
END TRAN
BEGIN CATCH
    IF (@CustomerComplete = 1) AND (XACT_STATE() = 1)
    BEGIN
        COMMIT TRAN
    END
    ELSE IF XACT_STATE() = 1
        ROLLBACK TRAN
END CATCH
```

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

References:

<http://stackoverflow.com/questions/11444923/stored-procedure-to-update-multiple-tables>

NEW QUESTION 130

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

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

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

Index Seek (NonClustered)

Scan a particular range of rows from a nonclustered index.

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

You need to resolve the performance issue.

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

- A. Yes
- B. No

Answer: A

Explanation:

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

NEW QUESTION 135

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 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 examine delays in executed threads, including errors with specific queries and batches.

- 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 ays .dm_os_memory_objeccs query.
- E. Create a sp_configure 'max server memory' query.

Answer: B

NEW QUESTION 140

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 70-762 Exam with Our Prep Materials Via below:

<https://www.certleader.com/70-762-dumps.html>