

70-762 Dumps

Developing SQL Databases (beta)

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



NEW QUESTION 1

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
```

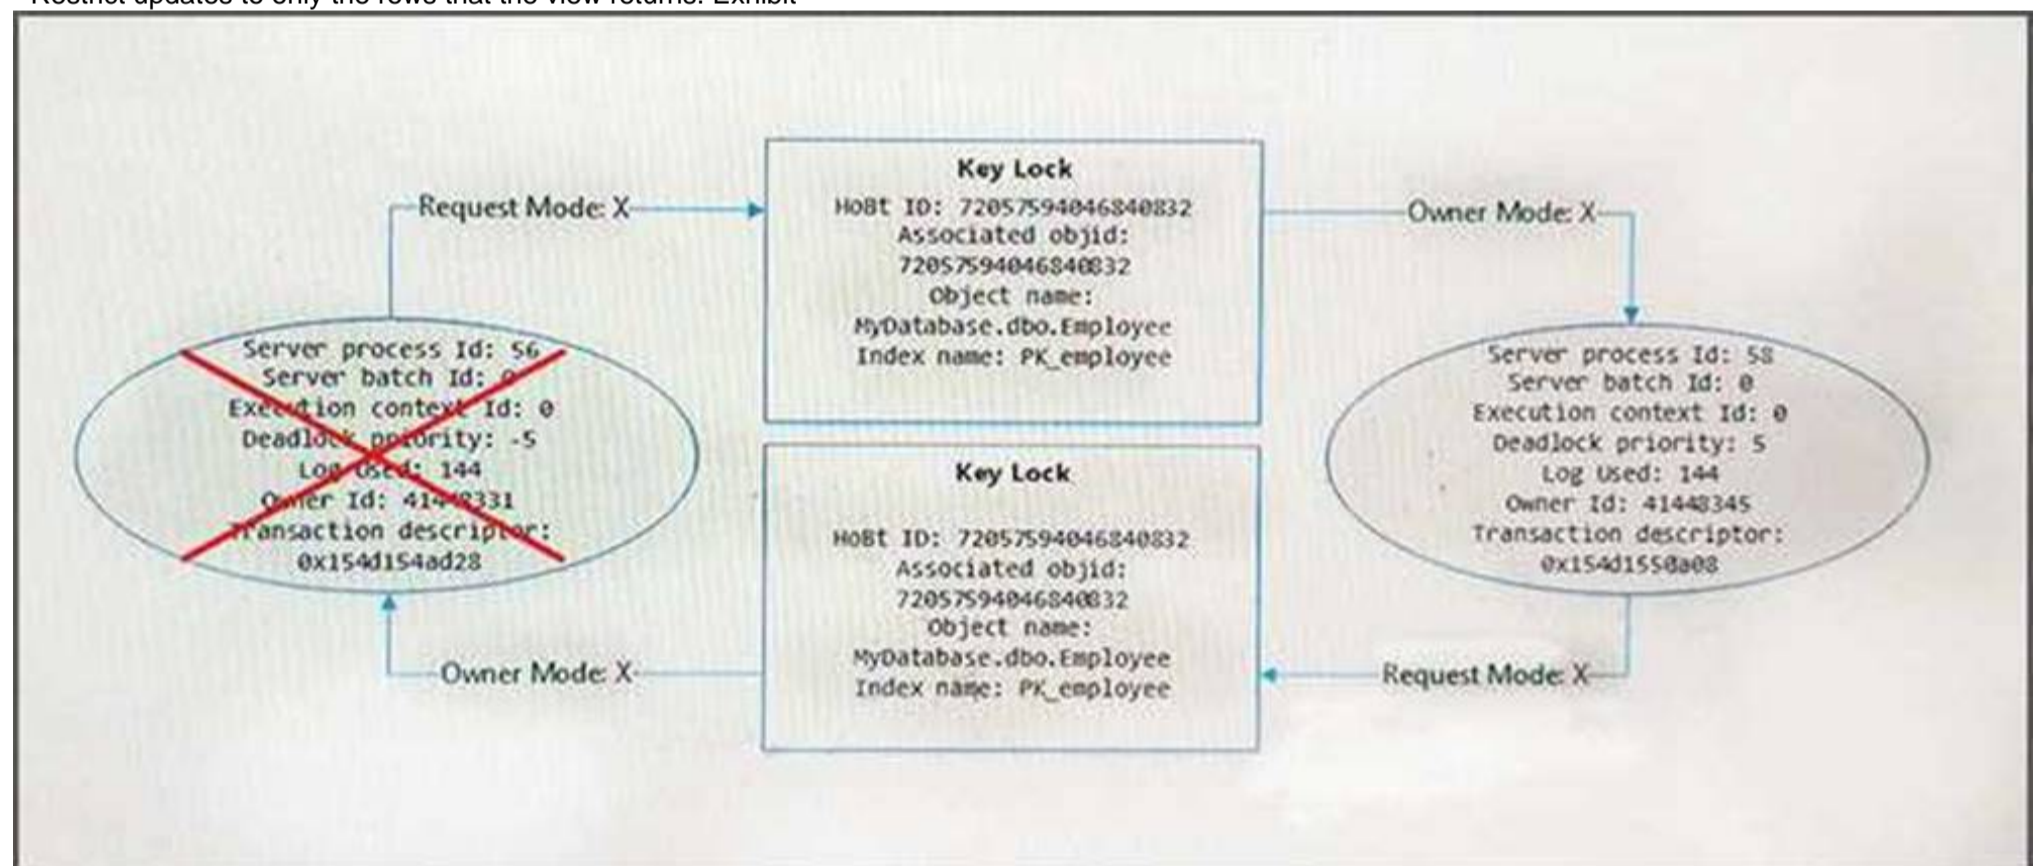
Application

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?

- Update the transaction level of the report query session to READPAST.
- Modify the report queries to use the UNION statement to combine the results of two or more queries.
- Set the READ_COMMITTED_SNAPSHOT database option to ON.
- Update the transaction level of the report query session to READ UNCOMMITTED.

Answer: C

Explanation:

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

NEW QUESTION 2

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 3

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

You are developing an application to track customer sales.

You need to create an object that meets the following requirements:

- Run managed code packaged in an assembly that was created in the Microsoft.NET Framework and uploaded in Microsoft SQL Server.
- Run within a transaction and roll back if a failure occurs.

- Run when a table is created or modified. What should you create?

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

Answer: B

Explanation:

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

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

NEW QUESTION 4

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

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

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

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

You frequently run the following Transact-SQL SELECT statements:

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

NEW QUESTION 5

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

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

You need to allow escalation of Table1 locks to the partition level and prevent all lock escalation for Table2. Which Transact-SQL statement should you run for each table? To answer, drag the appropriate Transact-SQL statements to the correct tables. Each command may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Transact-SQL statements

DBCC TRACEON(1211, -1)

DBCC TRACEON(1224, -1)

ALTER TABLE Table1 SET (LOCK_ESCALATION - DISABLE)

ALTER TABLE Table2 SET (LOCK_ESCALATION - DISABLE)

ALTER TABLE Table1 SET (LOCK_ESCALATION - AUTO)

ALTER TABLE Table2 SET (LOCK_ESCALATION - AUTO)

ALTER TABLE Table1 SET (LOCK_ESCALATION - TABLE)

ALTER TABLE Table2 SET (LOCK_ESCALATION - TABLE)

Answer Area

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

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

Box 1: Table1, Auto

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

Box 2: Table 2, DISABLE

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

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

References:

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

NEW QUESTION 6

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);

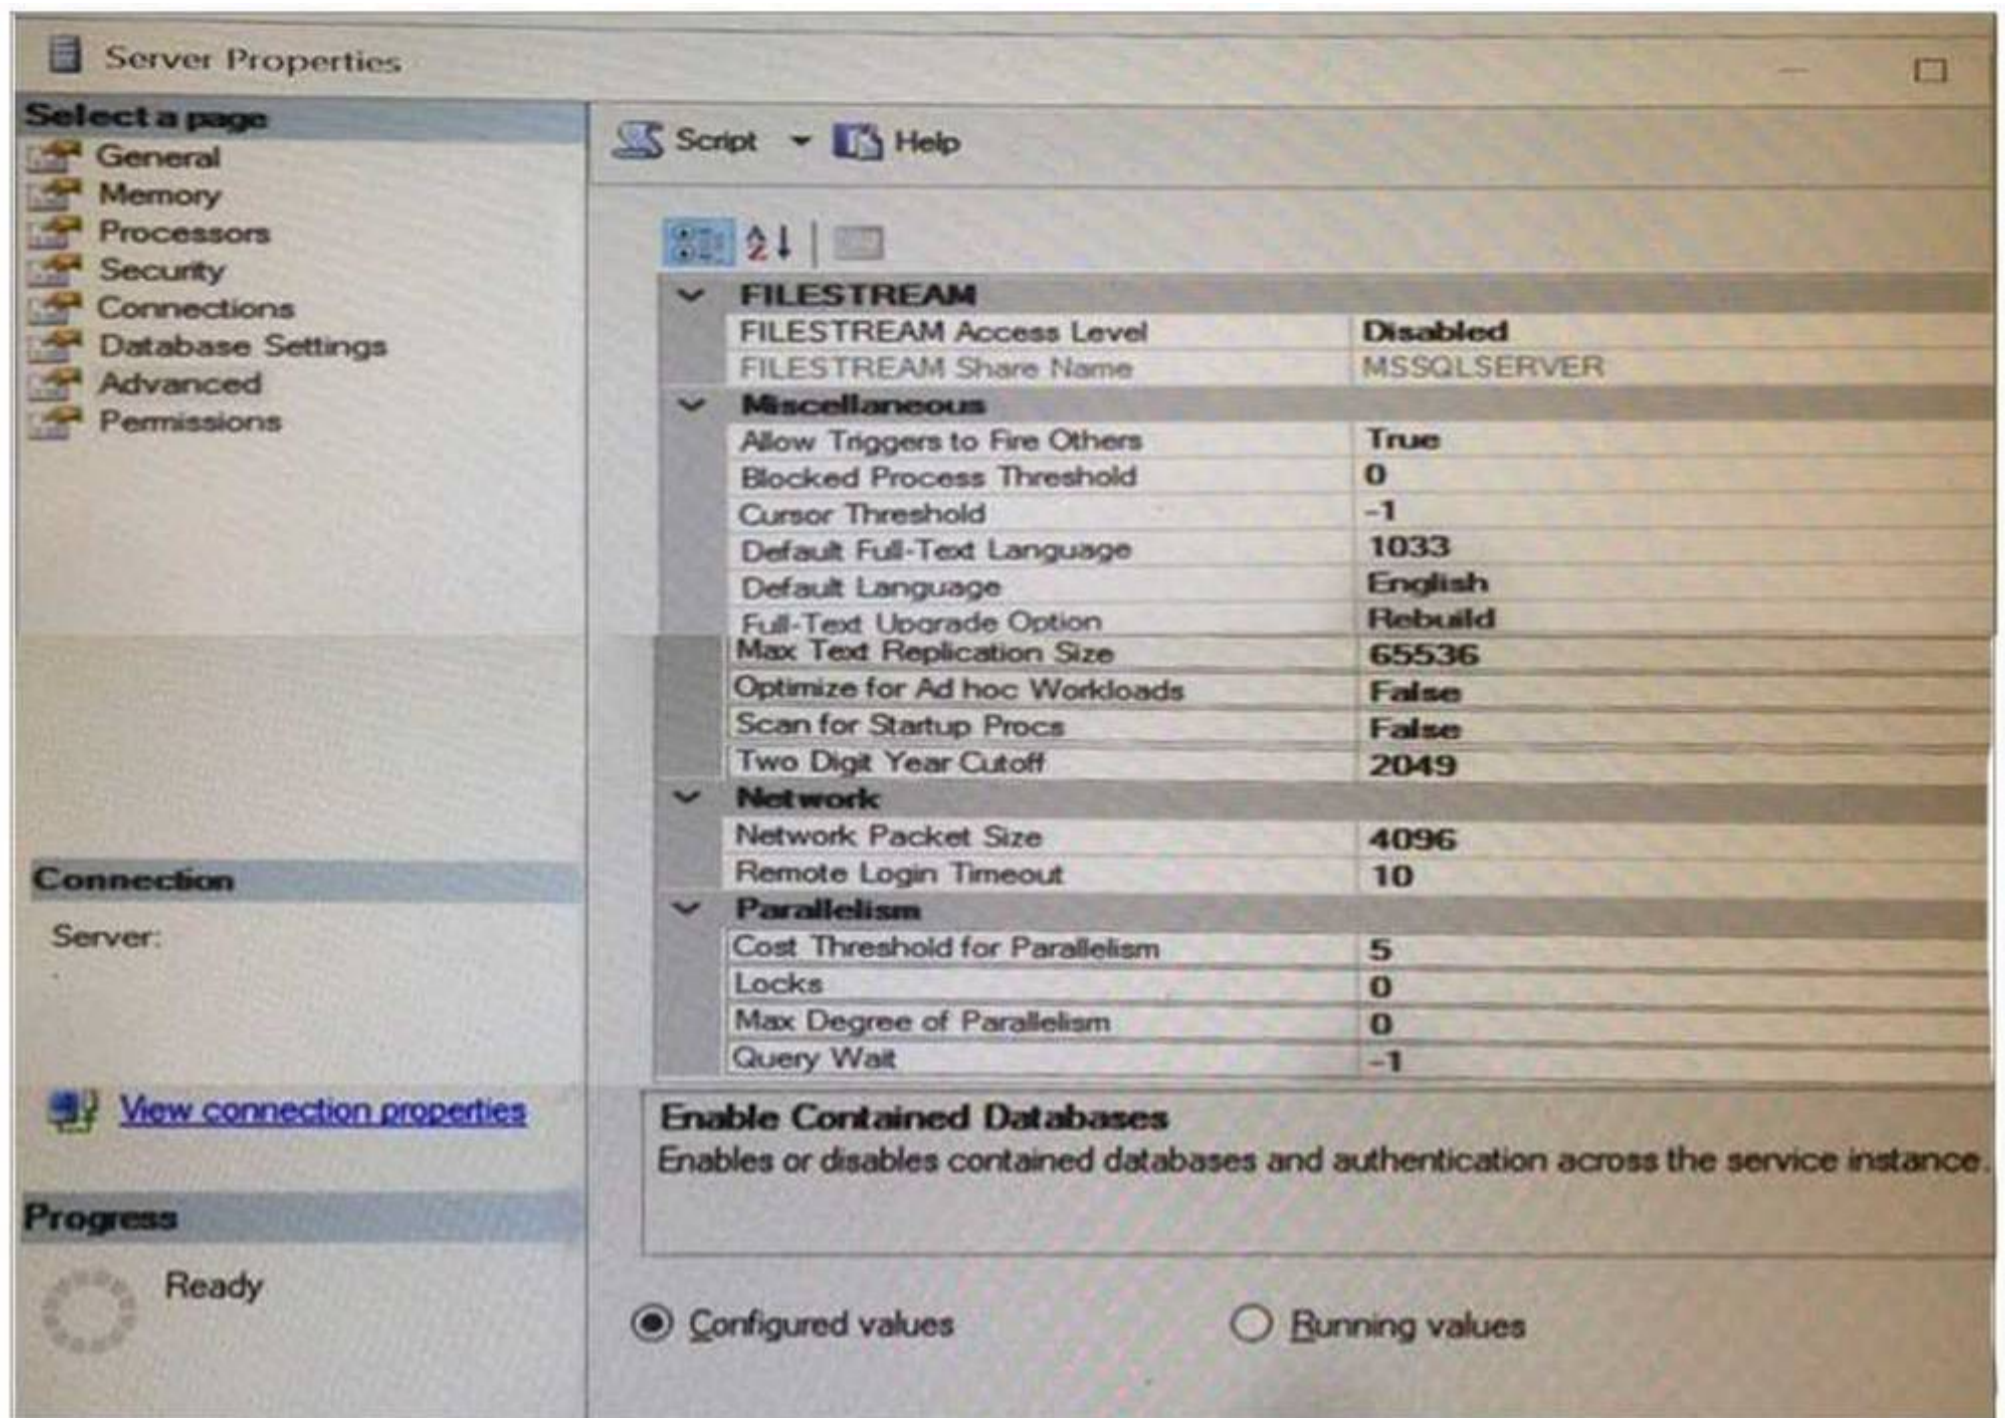
Answer: D

NEW QUESTION 7

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

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

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

* Optimize for ad hoc workloads

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

* Cost Threshold for Parallelism

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

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

NEW QUESTION 8

You have multiple stored procedures inside a transaction.

You need to ensure that all the data modified by the transaction is rolled back if a stored procedure causes a deadlock or times out.

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 XACT ABORT to ON.
- G. Execute the alter table T1 set (lock escalation = auto); statement.
- H. Use the output parameters.

Answer: B

NEW QUESTION 9

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 10

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution 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 10

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

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run during off-peak hours.

You observe that many deadlocks appear to be happening during specific times of the day.

You need to monitor the SQL environment and capture the information about the processes that are causing the deadlocks.

What should you do?

- A. Create a sys.dm_os_waiting_tasks query.
- B. Create a sys.dm_exec_sessions query.
- C. Create a PerformanceMonitor Data Collector Set.
- D. Create a sys.dm_os_memory_objects query.

- E. Create a sp_configure 'max server memory' query.
- F. Create a SQL Profiler trace.
- G. Create a sys.dm_os_wait_stats query.
- H. Create an Extended Event.

Answer: F

Explanation:

To view deadlock information, the Database Engine provides monitoring tools in the form of two trace flags, and the deadlock graph event in SQL Server Profiler. Trace Flag 1204 and Trace Flag 1222

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

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

NEW QUESTION 13

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

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run during off-peak hours.

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

You need to collect query performance data while minimizing the performance impact on the SQL Server. What should you do?

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

Answer: C

Explanation:

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

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

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

Predefined data collection sets include:

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

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

References:

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

NEW QUESTION 14

You have an existing Microsoft SQL Trace script. You plan to convert the script to an Extended Events session.

You need to collect the trace ID and other required information. Which system table should you use?

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

Answer: A

NEW QUESTION 17

You have a database that contains three encrypted store procedures named dbo.Proc1, dbo.Proc2 and dbo.Proc3. The stored procedures include INSERT, UPDATE, DELETE and BACKUP DATABASE statements.

You have the following requirements:

* You must run all the stored procedures within the same transaction.

* You must automatically start a transaction when stored procedures include DML statements.

* You must not automatically start a transaction when stored procedures include DDL statements. You need to run all three stored procedures.

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

Transact-SQL segments

BEGIN CATCH
IF (XACT_STATE() != 0)
ROLLBACK TRANSACTION
END CATCH

IF (@TRANCOUNT > 0)
ROLLBACK TRANSACTION

BEGIN TRAN

EXEC dbo.Proc1
EXEC dbo.Proc2
EXEC dbo.Proc3

SET IMPLICIT_TRANSACTIONS OFF

SET IMPLICIT_TRANSACTIONS ON

COMMIT TRANSACTION

BEGIN TRY
EXEC dbo.Proc1
EXEC dbo.Proc2
EXEC dbo.Proc3
IF (XACT_STATE() = 1)
COMMIT TRANSACTION;
END TRY

Answer Area



- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Note:

Implicit transaction mode remains in effect until the connection executes a SET IMPLICIT_TRANSACTIONS OFF statement, which returns the connection to autocommit mode. In autocommit mode, all individual statements are committed if they complete successfully.

When a connection is in implicit transaction mode and the connection is not currently in a transaction, executing any of the following statements starts a transaction:

Note 2: XACT_STATE returns the following values.

1 The current request has an active user transaction. The request can perform any actions, including writing data and committing the transaction. The transaction is committable.

-1 The current request has an active user transaction, but an error has occurred that has caused the transaction to be classified as an uncommittable transaction. The transaction is uncommittable and should be rolled back.

0 There is no active user transaction for the current request. A commit or rollback operation would generate an error.

References:

[https://technet.microsoft.com/en-us/library/ms187807\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms187807(v=sql.105).aspx) [https://technet.microsoft.com/en-us/library/ms189797\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/ms189797(v=sql.110).aspx)

NEW QUESTION 20

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

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

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

Transact-SQL segments

sys.dm_exec_query_stats o

sys.dm_db_partition_stats o

sys.dm_exec_sessions o

sys.dm_tran_database_transactions o

highest_cpu_queries.plan_handle DESC

highest_cpu_queries.total_worker_time DESC

q.objectid DESC

q.number DESC

Answer Area

SELECT
highest_cpu_queries.plan_handle,
highest_cpu_queries.total_worker_time,
q.dbid,
q.objectid,
q.number,
q.encrypted,
q.[text]
FROM
(SELECT TOP 50
o.plan_handle,
o.total_worker_time
FROM
Transact-SQL segment
ORDER BY o.total_worker_time desc) AS highest_cpu_queries
CROSS APPLY sys.dm_exec_sql_text(plan_handle) AS q
ORDER BY Transact-SQL segment ;

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: sys.dm_exec_query_stats

sys.dm_exec_query_stats returns aggregate performance statistics for cached query plans in SQL Server. Box 2: highest_cpu_queries.total_worker_time DESC
Sort on total_worker_time column

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

USE AdventureWorks2012; GO

SELECT TOP 5 query_stats.query_hash AS "Query Hash",

SUM(query_stats.total_worker_time) / SUM(query_stats.execution_count) AS "Avg CPU Time", MIN(query_stats.statement_text) AS "Statement Text"

FROM (SELECT QS.*,

SUBSTRING(ST.text, (QS.statement_start_offset/2) + 1,

((CASE statement_end_offset

WHEN -1 THEN DATALENGTH(ST.text)

ELSE QS.statement_end_offset END

- QS.statement_start_offset)/2) + 1) AS statement_text FROM sys.dm_exec_query_stats AS QS

CROSS APPLY sys.dm_exec_sql_text(QS.sql_handle)as ST) as query_stats GROUP BY query_stats.query_hash

ORDER BY 2 DESC;

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

NEW QUESTION 23

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.

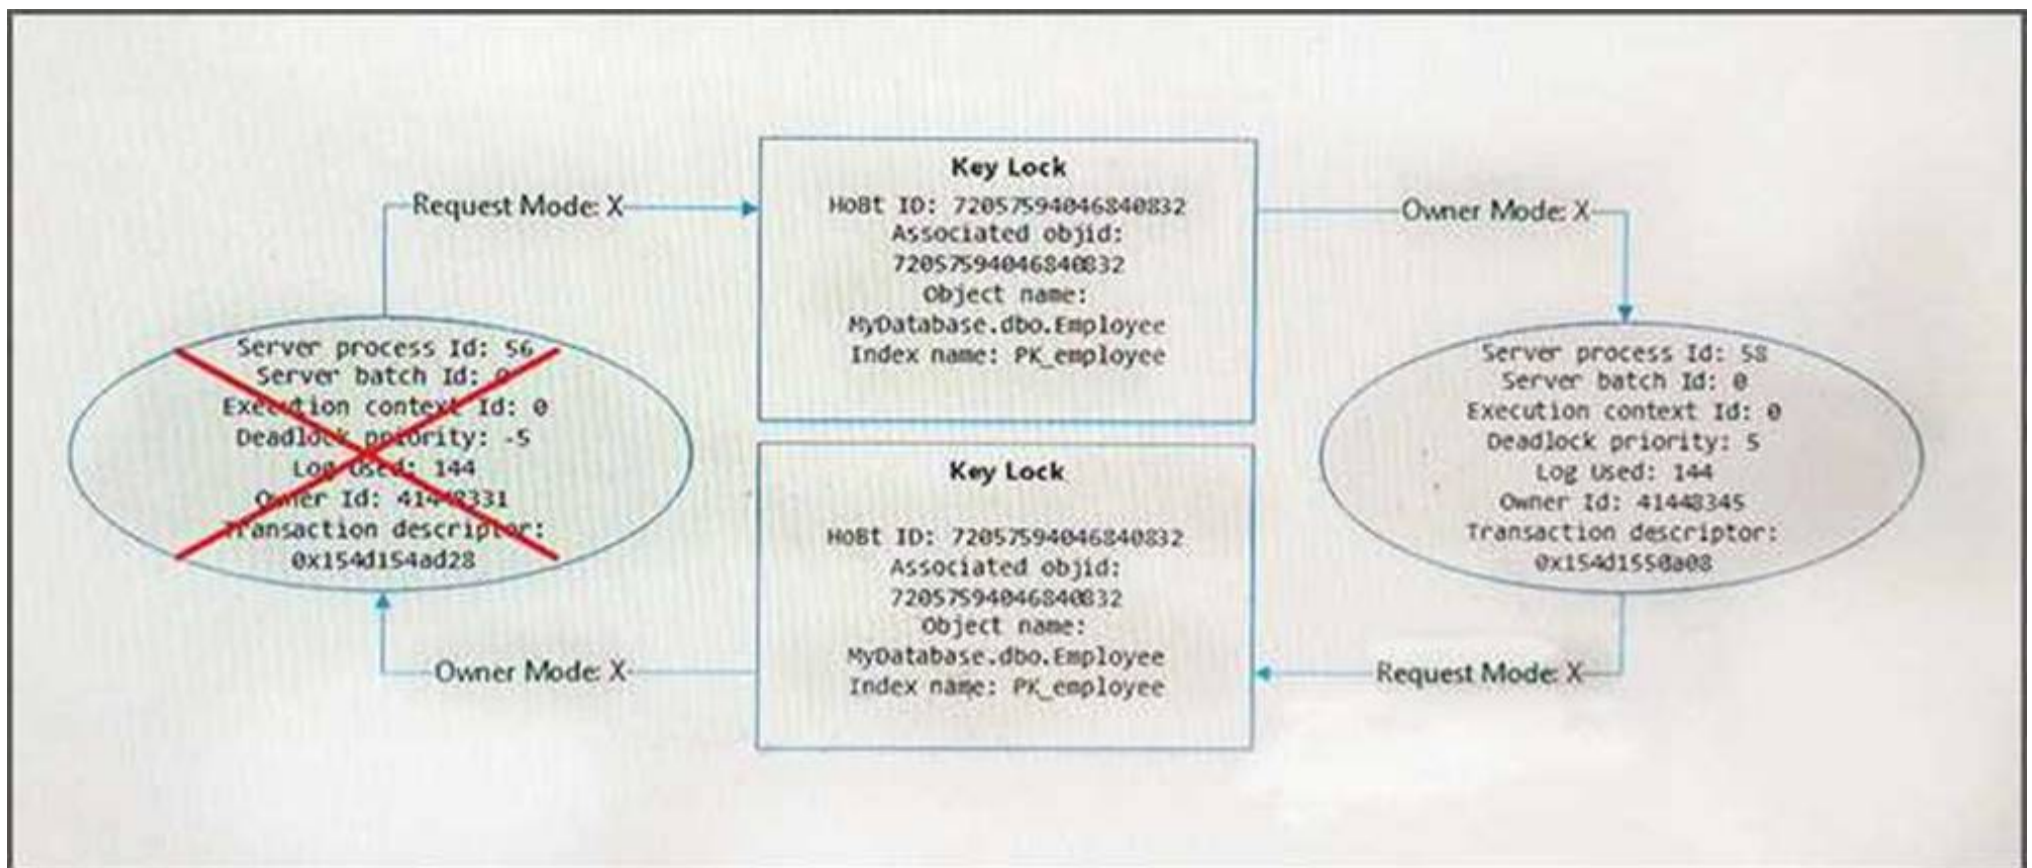
UspB: This stored procedure updates only the EmployeePayRate column.

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 28

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

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

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

Always return a value to the caller.

Return 0 if @Status is NULL.

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

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

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

Answer: BC

Explanation:

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

References:

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

NEW QUESTION 31

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

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

What should you add at line 02?

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

Answer: D

NEW QUESTION 33

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 NONCLUSTERED INDEX PK_Account ON Account(AccountNumber);
CREATE NONCLUSTERED INDEX IX_Account_ProductCode ON Account(Product-
Code) INCLUDE (Balance);
```

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Create a clustered index on the AccountNumber column as it is unique, not a non nonclustered one. References: <https://msdn.microsoft.com/en-us/library/ms190457.aspx>

NEW QUESTION 34

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

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

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

Commands

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

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

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

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

```
DROP TRIGGER CheckTriggerCreation ON  
DATABASE;
```

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

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Commands

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

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

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

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

```
DROP TRIGGER CheckTriggerCreation ON
DATABASE;
```

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

Answer Area

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

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

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



NEW QUESTION 37

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

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

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

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

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 39

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

Your company has employees in different regions around the world.

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

Syntax: datetimeoffset [(fractional seconds precision)]

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

NEW QUESTION 40

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

A new web dashboard is released to present the data from the sensors. Engineers report that the applications take longer than expected to commit updates.

You need to change the dashboard queries to improve concurrency and to support reading uncommitted data. What should you do?

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

Answer: C

NEW QUESTION 43

You 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 45

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(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 50

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

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

What should you do?

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

Answer: D

NEW QUESTION 53

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

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

You need to implement and enforce the following business rules:

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

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

- A. Yes
- B. No

Answer: A

Explanation:

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

NEW QUESTION 55

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

read committed

serializable

read uncommitted

repeatable read

Answer area

Scenario

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

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

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

Isolation levels

Isolation level

Isolation level

Isolation level

- 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 57

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 a sys.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 61

You have a Microsoft Azure SQL database. Some queries take a long time to execute.
You need to identify whether any of the long-runfiing queries prevent users from updating data in the database.
Which query should you execute? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```

SELECT
    t1.resource_type,
    t1.resource_associated_entity_id,
    t1.request_mode,
    t1.request_session_id,
    t2.blocking_session_id
FROM
    sys.dm_os_dispatchers as t1
    INNER JOIN sys.dm_tran_locks as t2
ON t1.lock_owner_address = t2.resource_address
WHERE t1.resource_database_id = db_id('db1')

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```

    t1.request_session_id,
    t2.blocking_session_id
FROM sys.dm_tran_locks as t1
    INNER JOIN sys.dm_tran_database_transactions as t2
ON t1.lock_owner_address = t2.resource_address
WHERE t1.resource_database_id = db_id('db1')

```

NEW QUESTION 66

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

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

- A. Activity monitor
- B. sys.dm_exec_connections
- C. Azure portal
- D. sys.rtuource_usage
- E. sys . dm_db_resource_stats

Answer: BC

NEW QUESTION 68

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

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

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

Answer Area

Module	Object type
Module1	<div>▼</div> <div> interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function </div>
Module2	<div>▼</div> <div> interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function </div>

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Module 1: Interpreted stored procedure

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

Natively compiled stored procedures are Transact-SQL stored procedures compiled to native code that access memory-optimized tables. Natively compiled stored procedures allow for efficient execution of the queries and business logic in the stored procedure.

References:

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

NEW QUESTION 72

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

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

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

You must optimize the indexes without making changes to the ix_FactSales_EmployeeKey index. You need to implement a columnstore index for the table. How should you complete the Transact-SQL code? To answer, select the appropriate Transact-SQL segments in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

▼

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

▼

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

▼

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

NEW QUESTION 77

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

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

You have a database that is 130 GB and contains 500 million rows of data.

Granular transactions and mass batch data imports change the database frequently throughout the day. Microsoft SQL Server Reporting Services (SSRS) uses the database to generate various reports by using several filters.

You discover that some reports time out before they complete. You need to reduce the likelihood that the reports will time out. Solution: You partition the largest tables.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

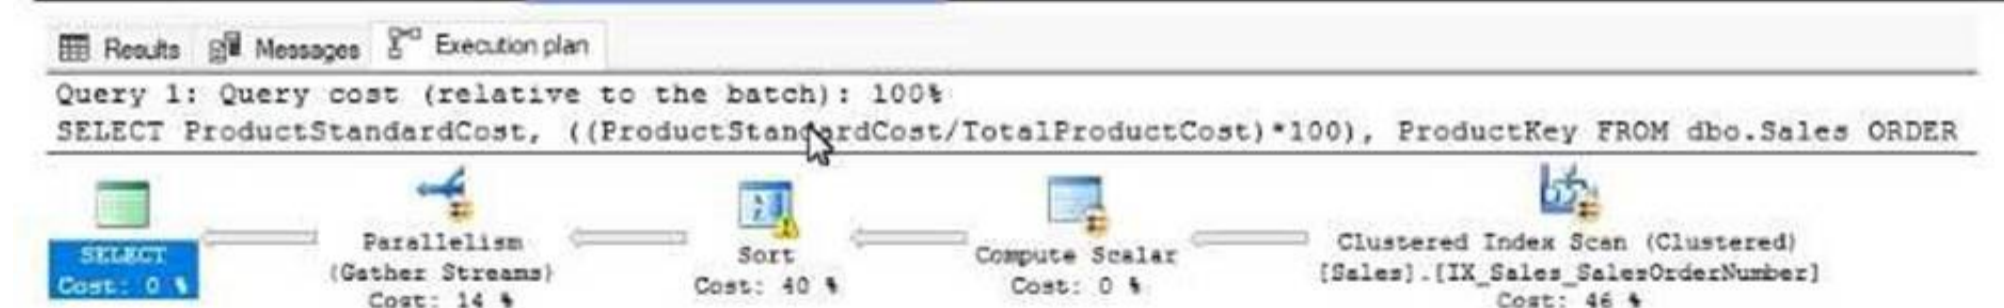
NEW QUESTION 79

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

You run the following Transact-SQL statement:

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

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



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

CREATE INDEX IX_Sales_01

ProductStandardCost DESC
ProductStandardCost ASC, ProductKey ASC
ProductStandardCost ASC, ProductKey DESC
ProductStandardCost DESC, ProductKey DESC

ON Sales (

WITH (ONLINE = ON)
INCLUDE (TotalProductCost)
WITH (SORT_IN_TEMPDB = ON)
WITH (SORT_IN_TEMPDB = OFF)

- A. Mastered
B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 82

You suspect deadlocks on a database.

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

NOTE: Each correct selection is worth one point.

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

Answer: AC

Explanation:

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

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

References:

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

NEW QUESTION 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 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

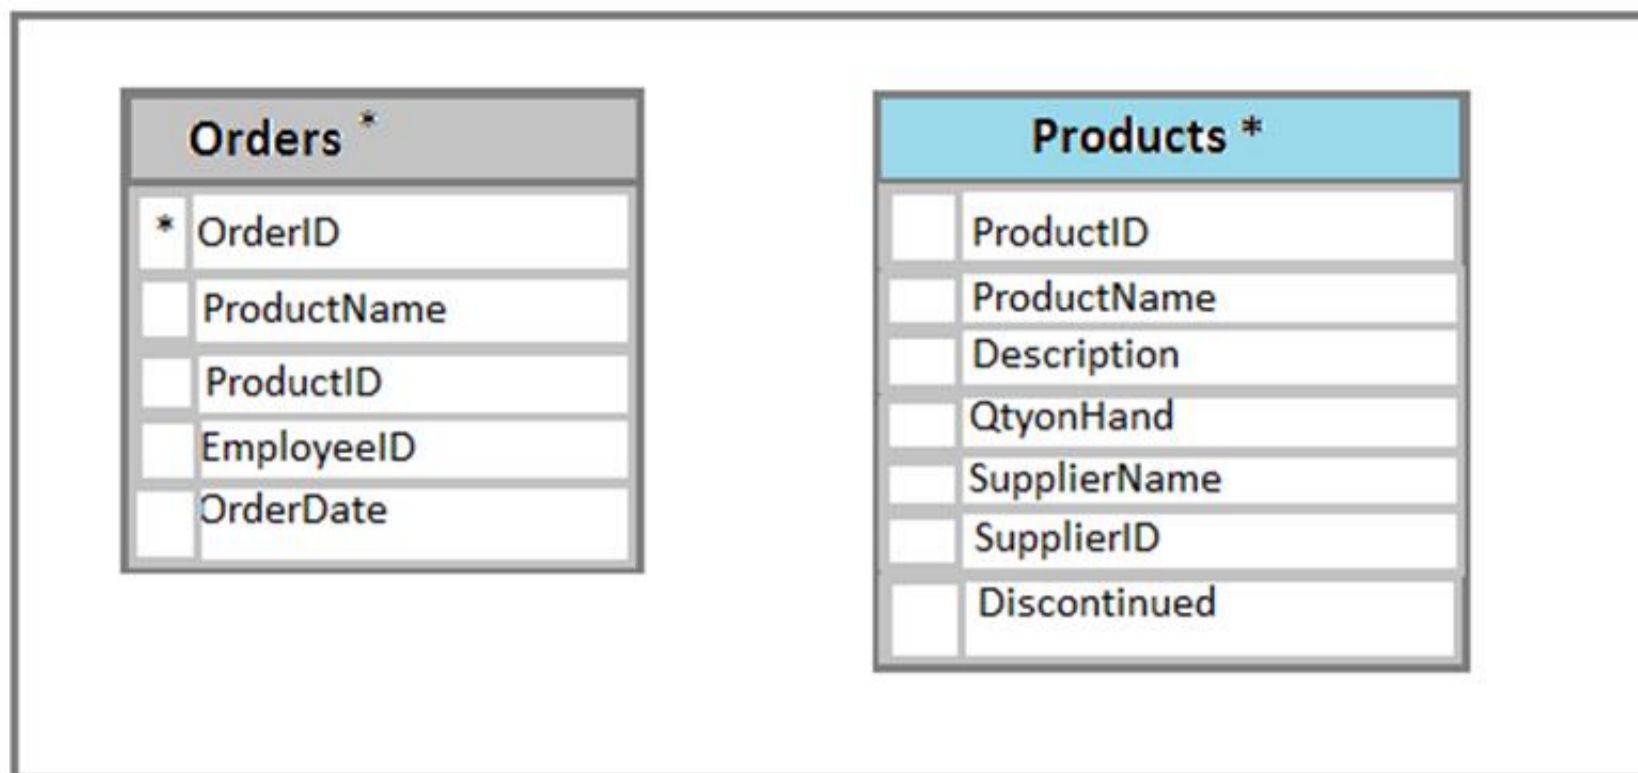
Explanation:

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

NEW QUESTION 88

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 91

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 96

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

You have a database named DB1. The database does not use a memory-optimized filegroup. The database contains a table named Table1. The table must support the following workloads:

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

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

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

Answer: C

Explanation:

A filtered index is an optimized nonclustered index, especially suited to cover queries that select from a well-defined subset of data. It uses a filter predicate to index a portion of rows in the table. A well-designed

filtered index can improve query performance, reduce index maintenance costs, and reduce index storage costs compared with full-table indexes.

References: [https://technet.microsoft.com/en-us/library/cc280372\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/cc280372(v=sql.105).aspx)

NEW QUESTION 98

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 101

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

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

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

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

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

- A. Yes
- B. No

Answer: A

Explanation:

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

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

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

NEW QUESTION 102

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

You are developing an application to track customer sales.

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 103

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 modify the design of the Orders table. What should you create?

A. a stored procedure with the RETURN statement

B. a FOR UPDATE trigger

C. an AFTER UPDATE trigger

D. a user defined function

Answer: D

Explanation:

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

1. Create new rows in the table without granting INSERT permissions to the table.

2. Notify the sales person who places an order whether or not the order was completed. References: <https://msdn.microsoft.com/en-us/library/ms186755.aspx>

NEW QUESTION 105

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

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

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

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

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

* Allow users access to the CustomerName and CustomerNumber columns for active customers.

* Allow changes to the columns that the view references. Modified data must be visible through the view.

* Prevent the view from being published as part of Microsoft SQL Server replication. Sales.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:

You run the uspUpdateCustomerSummary stored procedure to make changes to customer account summaries. Other stored procedures call the

spDeleteCustAcctRelationship to delete records from the CustomerToAccountBridge table.

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

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

A. The value of the TotalAccountCount column decreased.

B. The value of the TotalDepositAccountCount column is not changed.

C. The statement that modifies TotalDepositAccountCount is excluded from the transaction.

D. The value of the TotalAccountCount column is not changed.

Answer: D

NEW QUESTION 107

You have the following stored procedure:

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: B

NEW QUESTION 109

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

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

You need to improve the application response time.

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

- A. Yes
- B. No

Answer: B

Explanation:

References:

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

NEW QUESTION 113

You have two databases with the following settings:

Setting	Value
DELAYED_DURABILITY	ALLOWED
MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT	ON

You run the following Transact-SQL statements:

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

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

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

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

Transact-SQL statements

```
COMMIT TRANSACTION WITH (DELAYED_DURABILITY =  
OFF)
```

```
COMMIT TRANSACTION WITH (DELAYED_DURABILITY = ON)
```

```
IF OBJECT_ID ('tempdb. .#DiskTable') IS NOT NULL  
    DROP TABLE #Disktable  
SELECT * INTO #DiskTable from DiskDb.DiskTable
```

```
BEGIN TRANSACTION
```

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

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

```
ALTER DATABASE DiskDb SET DELAYED_DURABILITY =  
FORCED
```

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: BEGIN TRANSACTION

Box 2: UPDATE ... #Disktable

Box 3: IF... SELECT INTO ...#Disktable Box 4: .. DELAYED_DURABILITY = ON

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

Syntax:

```
COMMIT [ { TRAN | TRANSACTION } ] [ transaction_name | @tran_name_variable ] [ WITH ( DELAYED_DURABILITY = { OFF | ON } ) ]
```

References:

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

NEW QUESTION 116

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

You have a database that contains the following tables: BlogCategory, BlogEntry, ProductReview, Product, and SalesPerson. The tables were created using the following Transact SQL statements:

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

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

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

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

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

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

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

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

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

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

You must add the following constraints to the SalesHistory table:

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

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

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

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

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

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

Answer Area

CREATE FUNCTION ▼ Sales.YIDSalesByPerson

(@SalesPersonID int, @minYIDSales money)

RETURNS TABLE ▼

AS

```
RETURN (SELECT TOP(@SalesPersonID) BusinessEntityID, SalesYID
FROM Sales.SalesPerson
WHERE SalesYID > @minYIDSales
ORDER BY SalesYID desc);
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

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

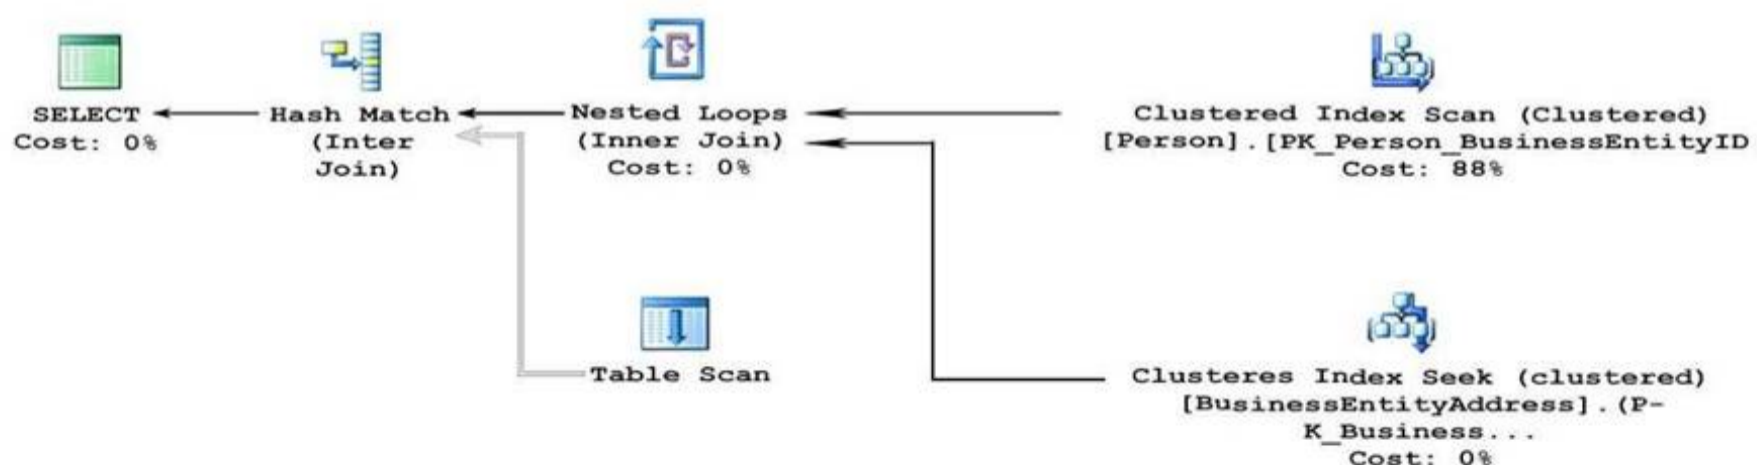
NEW QUESTION 121

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 columnstoreindex on the table.
- F. Create a hash index on the table.

Answer: B

Explanation:

By including nonkey columns, you can create nonclustered indexes that cover more queries. This is because the nonkeycolumns have the following benefits:
They can be data typesnot 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 125

You have a data warehouse fact table that has a clustered columnstore index. You have multiple CSV files that contain a total of 3 million rows of data. You need to upload the data to the fact table. The solution must avoid the delta group when you import the data. Which solution will achieve the goal in the least amount of time?

- A. Load the source data to a staging tabl
- B. Load the data to the fact table by using the insert_select statement and specify the Tablock option on the staging table.
- C. Load the source data to a staging table that has a clustered index on the primary ke
- D. Copy the data to the fact table by using the insert_select statement.
- E. Load the source data to the fact table by running bcp.exe and specify the _ Tablock option.
- F. Load the source data to the fact table by using the bulk insert statement and specify the Tablock option.

Answer: D

NEW QUESTION 126

You need to implement triggers to automate responses to the following events:

SQL Server logons
Database schema changes
Database updates

Which trigger types should you use? To answer, drag the appropriate trigger types to the appropriate scenarios. Each trigger type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Trigger types	Answer Area	Trigger type
LOGON	Scenario	
INSTEAD OF INSERT	Prevent a user from establishing a session if they have an existing session	Trigger
AFTER INSERT	Insert into tables when data is inserted into a non-updatable view	Trigger
DDL	Insert the column name, table name, and user name when a column is added to a table	Trigger

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: LOGON

Logon triggers fire stored procedures in response to a LOGON event. This event is raised when a user session is established with an instance of SQL Server.

Box 2: INSTEAD OF INSERT

An "INSTEAD of trigger" is executed instead of the original operation, and not combining with the operation. INSTEAD OF triggers override the standard actions of the triggering statement. It can be used to bypass the statement and execute a whole different statement, or just help us check and examine the data before the action is done.

Box 3: DDL

DDL statements (CREATE or ALTER primarily) issued either by a particular schema/user or by any schema/user in the database

Note:

You can write triggers that fire whenever one of the following operations occurs:

DML statements (INSERT, UPDATE, DELETE) on a particular table or view, issued by any user

DDL statements (CREATE or ALTER primarily) issued either by a particular schema/user or by any schema/user in the database

Database events, such as logon/logoff, errors, or startup/shutdown, also issued either by a particular schema/user or by any schema/user in the database

References: https://docs.oracle.com/cd/B19306_01/server.102/b14220/triggers.htm

<https://social.technet.microsoft.com/wiki/contents/articles/28152.t-sql-instead-of-triggers.aspx> <https://docs.microsoft.com/en-us/sql/relational-databases/triggers/logon-triggers?view=sql-server-2017>

NEW QUESTION 129

You are analyzing the performance of a database environment. You need to find all unused indexes in the current database.

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

Answer Area

```

SELECT a.name as TableName
       b.name as IndexName

FROM   sys.indexes b

INNER JOIN sys.dm_exec_query_stats s ON s.object_id = b.object_id AND s.index_id = b.index_id

INNER JOIN sys.indexes a ON b.object_id = a.object_id

WHERE ((user_seeks = 0 AND user_scans = 0 AND user_lookups = 0) OR s.object_id is null)
  
```

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Example: Following query helps you to find all unused indexes within database using sys.dm_db_index_usage_stats DMV.

-- Ensure a USE statement has been executed first. SELECT u.*

FROM [sys].[indexes] i

INNER JOIN [sys].[objects] o ON (i.OBJECT_ID = o.OBJECT_ID)

LEFT JOIN [sys].[dm_db_index_usage_stats] u ON (i.OBJECT_ID = u.OBJECT_ID) AND i.[index_id] = u.[index_id]

AND u.[database_id] = DB_ID() --returning the database ID of the current database WHERE o.[type] <>'S' --shouldn't be a system base table

AND i.[type_desc] <> 'HEAP' AND i.[name] NOT LIKE 'PK_ %'

AND u.[user_seeks] + u.[user_scans] + u.[user_lookups] = 0 AND u.[last_system_scan] IS NOT NULL

ORDER BY 1 ASC

References: https://basitaalishan.com/2012/06/15/find-unused-indexes-using-sys-dm_db_index_usage_stats/

NEW QUESTION 131

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

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run during off-peak hours.

You must monitor and optimize the SQL Server to maximize throughput, response time, and overall SQL performance.

You need to identify previous situations where a modification has prevented queries from selecting data in tables.

What should you do?

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

Answer: G

Explanation:

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

NEW QUESTION 133

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 following tables:

Table name	Description
TBL1	<ul style="list-style-type: none"> Column1 is configured as the primary key. Column2 will store the year. TBL1 only stores data for the year 2016. The table will contain 1 million records.
TBL2	<ul style="list-style-type: none"> Column1 is configured as the primary key. Column2 will store the year. TBL2 only stores data for the year 2015. The table will contain 1 million records.

Users frequently run the following query. The users report that the query takes a long time to return results.

```
SELECT Column1, Column2, Column3
FROM (
    SELECT Column1, Column2, Column3
    FROM TBL1
    UNION ALL
    SELECT Column1, Column2, Column3
    FROM TBL2)
WHERE Column2 = <year> AND Column3 = 1
```

You need to minimize the amount of time required for the query 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 a 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.
- I. Make no changes to TBL2.
- J. Create CHECK constraints on both TBL1 and TBL2. Create a partitioned view that combines columns from TBL1 and TBL2.
- K. Create an indexed view that combines columns from TBL1 and TBL2.

Answer: H

NEW QUESTION 138

You have a database that contains both disk-based and memory-optimized tables.

You need to create two modules. The modules must meet the requirements described in the following table.

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

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

Answer Area

Module	Object type
Module 1	<div> ▼ <ul style="list-style-type: none"> interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function </div>
Module 2	<div> ▼ <ul style="list-style-type: none"> interpreted stored procedure multi-statement table-valued function natively compiled stored procedure natively compiled user-defined function user-defined scalar function </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Returning Data by Using OUTPUT Parameters

If you specify the OUTPUT keyword for a parameter in the procedure definition, the stored procedure can return the current value of the parameter to the calling program when the stored procedure exits.

SQL Server stored procedures, views and functions are able to use the WITH ENCRYPTION option to disguise the contents of a particular procedure or function from discovery.

Native Compilation of Tables and Stored Procedures

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

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

NEW QUESTION 141

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. DDL 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 142

You are developing queries and stored procedures to support a line-of-business application. You need to use the appropriate isolation level based on the scenario. Which isolation levels should you implement? To answer, select the appropriate isolation level for each scenario in the answer area. Each isolation level may be used only once.

NOTE: Each correct selection is worth one point.

Answer Area

Scenario

If a transaction is not completed then a dirty read is performed.

Transaction cannot read data that has been modified but not committed by other transactions.

Transaction may not read data that has been modified but not committed by other transactions. Other transactions may not modify data until the current transaction completes.

Uses row versioning to present each statement with a snapshot of the data as it existed at the start of the statement.

Isolation level

	▼
Read Uncommitted	
Read Committed	
Repeatable Read	
Snapshot	

	▼
Read Uncommitted	
Read Committed	
Repeatable Read	
Snapshot	

	▼
Read Uncommitted	
Read Committed	
Repeatable Read	
Snapshot	

	▼
Read Uncommitted	
Read Committed	
Repeatable Read	
Snapshot	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: READ UNCOMMITTED

Transactions running at the READ UNCOMMITTED level do not issue shared locks to prevent other transactions from modifying data read by the current transaction. READ UNCOMMITTED transactions are also not blocked by exclusive locks that would prevent the current transaction from reading rows that have been modified but not committed by other transactions. When this option is set, it is possible to read uncommitted modifications, which are called dirty reads. Values in the data can be changed and rows can appear or disappear in the data set before the end of the transaction.

Box 2: READ COMMITTED

READ COMMITTED specifies that statements cannot read data that has been modified but not committed by other transactions.

SERIALIZABLE specifies that statements cannot read data that has been modified but not yet committed by other transactions.

Box 3: REPEATABLE READ

REPEATABLE READ specifies that statements cannot read data that has been modified but not yet committed by other transactions and that no other transactions can modify data that has been read by the current transaction until the current transaction completes.

Box 4: SNAPSHOT

If READ_COMMITTED_SNAPSHOT is set to ON, the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement.

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-transaction-isolation-level-transact-sql?view=sql-server>

NEW QUESTION 146

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.

Orders *	
<input checked="" type="checkbox"/>	OrderID
<input type="checkbox"/>	ProductName
<input type="checkbox"/>	ProductID
<input type="checkbox"/>	EmployeeID
<input type="checkbox"/>	OrderDate

Products *	
<input type="checkbox"/>	ProductID
<input type="checkbox"/>	ProductName
<input type="checkbox"/>	Description
<input type="checkbox"/>	QtyonHand
<input type="checkbox"/>	SupplierName
<input type="checkbox"/>	SupplierID
<input type="checkbox"/>	Discontinued

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

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

You need to create triggers that meet the following requirements:

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

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

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

Answer Area

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

INSTEAD OF INSERT triggers can be defined on a view or table to replace the standard action of the INSERT statement.

AFTER specifies that the DML trigger is fired only when all operations specified in the triggering SQL statement have executed successfully.

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

NEW QUESTION 148

You are performing a code review for Transact-SQL statements.

What action does the constraint in each statement perform? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

• • • • •

Answer Area

SQL Statement

```
CREATE TABLE Persons (
    ID INT NOT NULL,
    LastName VARCHAR(255) NOT NULL,
    FirstName VARCHAR(255),
    Age INT,
    CONSTRAINT PK_Person PRIMARY KEY (ID, LastName)
)

CREATE TABLE Orders (
    OrderID INT NOT NULL,
    OrderNumber INT NOT NULL,
    PersonID INT,
    PRIMARY KEY (OrderID),
    CONSTRAINT FK_PersonOrder FOREIGN KEY (PersonID)
    REFERENCES Persons(PersonID)
)

CREATE TABLE Persons (
    ID INT NOT NULL UNIQUE,
    LastName VARCHAR(255) NOT NULL,
    FirstName VARCHAR(255),
    Age INT
)

CREATE TABLE Persons (
    ID INT NOT NULL,
    LastName VARCHAR(255) NOT NULL,
    FirstName VARCHAR(255),
    Age INT CHECK (Age>=18)
)
```

Action

Enforce unique values for columns.
Enforce a reference to a primary key.
Ensure unique values within a column.
Ensure values are acceptable for columns based on an evaluation.

Enforce unique values for columns.
Enforce a reference to a primary key.
Ensure unique values within columns.
Ensure values are acceptable for columns based on an evaluation.

Enforce unique values for a columns.
Enforce a reference to a primary key.
Ensure unique values within columns.
Ensure values are acceptable for columns based on an evaluation.

Enforce unique values for a columns.
Enforce a reference to a primary key.
Ensure unique values within columns.
Ensure values are acceptable for columns based on an evaluation.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SQL Statement

Action

```
CREATE TABLE Persons (
    ID INT NOT NULL,
    LastName VARCHAR(255) NOT NULL,
    FirstName VARCHAR(255),
    Age INT,
    CONSTRAINT PK_Person PRIMARY KEY (ID, LastName)
)

CREATE TABLE Orders (
    OrderID INT NOT NULL,
    OrderNumber INT NOT NULL,
    PersonID INT,
    PRIMARY KEY (OrderID),
    CONSTRAINT FK_PersonOrder FOREIGN KEY (PersonID)
    REFERENCES Persons(PersonID)
)

CREATE TABLE Persons (
    ID INT NOT NULL UNIQUE,
    LastName VARCHAR(255) NOT NULL,
    FirstName VARCHAR(255),
    Age INT
)

CREATE TABLE Persons (
    ID INT NOT NULL,
    LastName VARCHAR(255) NOT NULL,
    FirstName VARCHAR(255),
    Age INT CHECK (Age>=18)
)
```

Ensure unique values within a column.

Enforce a reference to a primary key.

Enforce unique values for a columns.

Ensure values are acceptable for columns based on an evaluation.

NEW QUESTION 151

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 following tables:

Table Name	Description
TBL1	<ul style="list-style-type: none"> The table will contain 10 million records. The frequency of inserting, updating, and deleting records is low.
TBL2	<ul style="list-style-type: none"> The table will contain 1 million records.

Users frequently run the following query:

```
SELECT TBL1.Column2, TBL2.Column2, SUM(TBL1.Column3), SUM(TBL1.Column4 * TBL1.Column5 * (TBL2.Column4 - TBL2.Column3))
FROM TBL1
INNER JOIN TBL2 TBL1.Column1 = TBL2.Column1
GROUP BY TBL1.Column2, TBL2.Column2
```

Users report that the query takes a long time to return results.

You need to minimize the amount of time required for the query to return data. What should you do?

- Create clustered indexes on TBL1 and TBL2.
- Create a clustered index on TBL1. Create a nonclustered index on TBL2 and add the most frequently queried columns as included columns.
- Create a nonclustered index on TBL2 only.
- Create unique constraints on both TBL1 and TBL2. Create a partitioned view that combines columns from TBL1 and TBL2.
- Drop existing indexes on TBL1 and then create a clustered columnstore index.
- Create a nonclustered columnstore index on TBL1. Create a nonclustered index on TBL2.
- Drop existing indexes on TBL1 and then create a clustered columnstore index.
- Create a nonclustered columnstore index on TBL1.
- Make no changes to TBL2.
- Create CHECK constraints on both TBL1 and TBL2. Create a partitioned view that combines columns from TBL1 and TBL2.
- Create an indexed view that combines columns from TBL1 and TBL2.

Answer: D

NEW QUESTION 154

You maintain a Microsoft Azure SQL Database instance.

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

details for each connection to the database

a list of all active user connections and internal tasks

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

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

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

Transact-SQL segments

GRANT SELECT ON SCHEMA: :sys TO User1

GRANT VIEW DATABASE STATE TO User1

GRANT VIEW DEFINITION TO User1

SELECT = FROM sys.dm_exec_connections

SELECT = FROM sys.dm_exec_requests

SELECT = FROM sys.dm_exec_sessions

Answer Area

CREATE PROCEDURE ViewConnections

AS

Transact-SQL segment

Transact-SQL segment

GO

Transact-SQL segment

- Mastered
- Not Mastered

Answer: A

Explanation:

Box 1: Sys.dm_exec_connections

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

Box 2: sys.dm_exec_sessions

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

Box 3: GRANT VIEW DATABASE STATE To User1

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

NEW QUESTION 155

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 158

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 159

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

*Returns multiple rows

*Optimizes the performance of the query within the function

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

Transact-SQL segments

RETURNS TABLE AS

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

RETURNS @Prices TABLE (ListPrice money) AS

RETURNS [Money] AS

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

Answer area

CREATE FUNCTION dbo.GetPrice (@Product

Transact-SQL segment

Transact-SQL segment

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

RETURNS @Prices TABLE (ListPrice money) AS

Answer area

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

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

NEW QUESTION 163

.....

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>