


Microsoft MCSA Certification 70-762 Exam



- **Vendor: Microsoft**
- **Exam Code: 70-762**
- **Exam Name: Developing SQL Databases**

Get Complete Version Exam 70-762 Dumps with VCE and PDF Here



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

QUESTION 1

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:

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

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:

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.

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 Sale Price 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: A

QUESTION 2

You are developing an application to track customer sales. You need to create an object that meet 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

Answer: C

QUESTION 3

Hotspot Question

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:

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

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:

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.

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 Sale Price 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 update the SalesHistory table. How should you complete the Transact_SQL statement? To answer? select the appropriate TransactSQL, segments in the answer area.

```
IF OBJECT_id(*SalesHistory*)>0 DROP TABLE SalesHistory
GO
IF OBJECT_ID(*ProductTypes*)>0 DROP TABLE ProductTypes
GO
CREATE TABLE ProductTypes
(
    ProductID SMALLINT,
    ProductDescription VARCHAR(255),
    CONSTRAINT pk_ProductID PRIMARY KEY (ProductID)
)
GO
```

```
CREATE TABLE [dbp].[SalesHistoryK]
[SalesID] [int]
```

IDENTITY(1,4)
IDENTITY(1,4) NOT NULL PRIMARY KEY
UNIQUE

```
[ProductID] SMALLINT NULL ,
```

```
[SaleDate] [datetime] NULL
```

```
[SalePrice] [money]
```

NOT NULL
NULL CHECK (SalesPrice > 4)
UNIQUE

```
[CategoryID] [smallint]
```

NOT NULL
NULL CHECK (SalesPrice > 4)
UNIQUE

CONSTRAINT fk_SalesHistoryProductID FOREIGN KEY (ProductID) REFERENCES SalesHistory(CategoryID)
CONSTRAINT fk_SalesHistoryProductID FOREIGN KEY (ProductID) REFERENCES ProductTypes(ProductID)

```
)
GO
```

Answer:

```
CREATE TABLE [dbp].[SalesHistoryK]
[SalesID] [int]
```

IDENTITY(1,4)
IDENTITY(1,4) NOT NULL PRIMARY KEY
UNIQUE

```
[ProductID] SMALLINT NULL ,
```

```
[SaleDate] [datetime] NULL
```

```
[SalePrice] [money]
```

NOT NULL
NULL CHECK (SalesPrice > 4)
UNIQUE

```
[CategoryID] [smallint]
```

NOT NULL
NULL CHECK (SalesPrice > 4)
UNIQUE

CONSTRAINT fk_SalesHistoryProductID FOREIGN KEY (ProductID) REFERENCES SalesHistory(CategoryID)
CONSTRAINT fk_SalesHistoryProductID FOREIGN KEY (ProductID) REFERENCES ProductTypes(ProductID)

QUESTION 4

Drag and Drop Question

You are analyzing the performance of a database environment. You suspect there are several missing indexes in the current database. You need to return a prioritized list of the missing indexes on the current database. How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct location s. Each Transact-SQL segment may be used once, more than once or not at all. You may need to drag the split bar between panes or scroll to view content.

Transact-SQL segments

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

Answer Area

```
SELECT so.name
      , (avg_total_user_cost * avg_user_impact) * (user_seeks + user_scans)
      as Impact
      , mid.equality_columns
      , mid.inequality_columns
      , mid.included_columns
FROM   Transact-SQL statement AS stats
INNER JOIN sys.dm_db_missing_index_groups AS mig ON stats.group_handle =
mig.index_group_handle
INNER JOIN sys.dm_db_missing_index_details AS mid ON mig.index_handle =
mid.index_handle
INNER JOIN sys.objects so WITH (nolock) ON mid.object_id = so.object_id
WHERE stats.group_handle IN (
SELECT TOP (5000) Transact-SQL statement
FROM   Transact-SQL statement WITH (nolocks)
ORDER BY (avg_total_user_cost * avg_user_impact) * (user_seeks +
user_scans DESC)
```

Answer:


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

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

SELECT TOP (5000) group_handle

FROM sys.dm_db_missing_index_group_stats WITH (nolocks)
ORDER BY (avg_total_user_cost * avg_user_impact) * (user_seeks +
user_scans DESC)
```

QUESTION 5**Drag and Drop Question**

You are monitoring a Microsoft Azure SQL Database. The database is experiencing high CPU consumption. You need to determine which query uses the most cumulative CPU. How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment may be used once, more than one or not at all. You may need to drag the split bar between panes or scroll to view content.

Transact-SQL segments

Answer Area

sys.dm_exec_query_stats o
sys.dm_db_partition_stats o
sys.dm_exec_sessions o
sys.dm_tran_database_transactions o
highest_cpu_queries.plan_handle DESC
highest_cpu_queries.total_worker_time DESC
q.objectid DESC
q.number DESC

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

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

ORDER BY Transact-SQL segment ;
```

Answer:

```
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
        sys.dm_exec_query_stats o

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

ORDER BY highest_cpu_queries.total_worker_time DESC
```

QUESTION 6

Drag and Drop Question

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

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

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

Transact-SQL segments

Answer Area

Requirement	Performance counter
Requirement 1	Performance counter
Requirement 2	Performance counter
Requirement 3	Performance counter

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

Answer:

Requirement	Performance counter
Requirement 1	SQL Server: Memory Manager: Total Server Memory (KB)
Requirement 2	SQL Server: Memory Manager: Granted Workspace Memory (KB)
Requirement 3	Process: working Set

QUESTION 7

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?

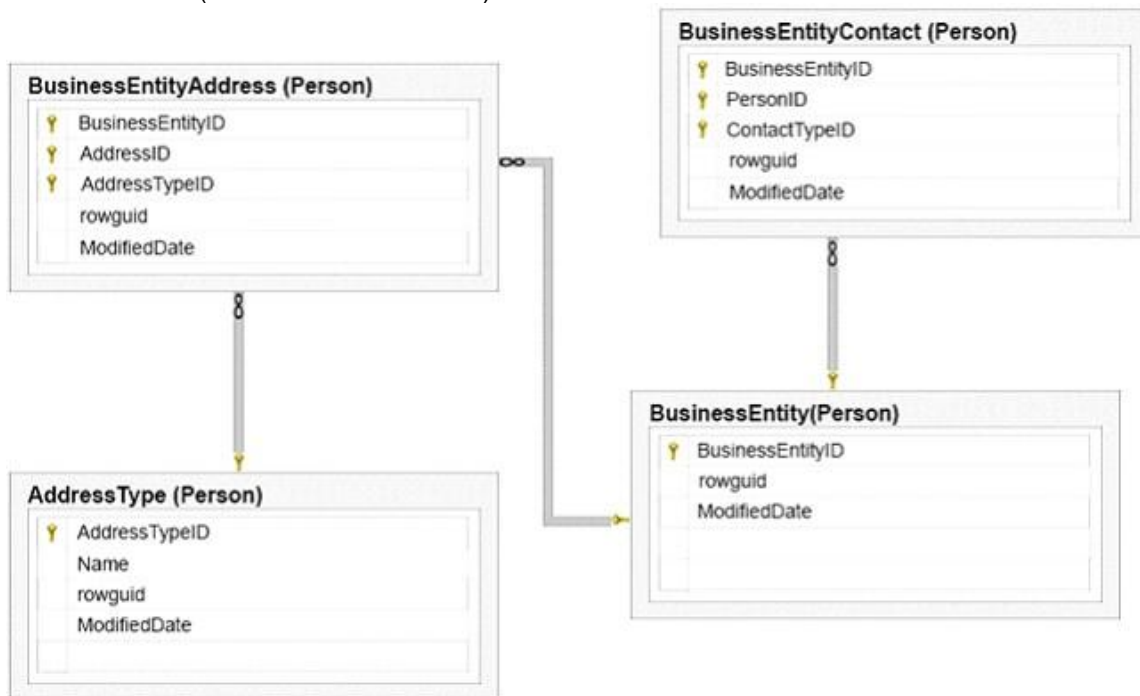
- A. table-valued function
- B. a schema-bound view
- C. a partitioned view
- D. a DML trigger

Answer: B

QUESTION 8

Drag and Drop Question

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



You need to insert a new customer record into the tables as a single unit of work. Which five Transact-SQL segments should you use to develop the SQLtution? To answer, move the appropriate Transact-SQL segments to the answer area and arrange the, in the correct order. NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Transact-SQL segments

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

Answer Area



Answer:

Transact-SQL segments

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

Answer Area



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

QUESTION 9

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

QUESTION 10

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. 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

QUESTION 11

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run during off-peak hours. You must monitor the SQL Server instances in real time and optimize the server to maximize throughput, response time, and overall SQL performance. What should you do?

- A. Create 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: B

QUESTION 12

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run during off-peak hours. You must monitor the SQL Server instances in real time and optimize the server to maximize throughput, response time, and overall SQL performance. You need to ensure that the performance of each instance is consistent for the same queried and query plans. What should you do?

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

Answer: H

QUESTION 13

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

QUESTION 14

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

Answer: D

QUESTION 15

Hotspot Question

You have a database named Sales. You need to create a table named Customer that includes the columns described in the following table.

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

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

Answer Area

CREATE TABLE Customer

(

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

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

MobileNO varchar(10)

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

NULL,

Email varchar(100)

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

NULL,

);

Answer:

Answer Area

```
CREATE TABLE Customer
```

```
(
```

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

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

```
    MobileNO varchar(10)
```

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

NULL,

```
    Email varchar(100)
```

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

NULL,

QUESTION 16

Hotspot Question

You are developing an app that allows users to query historical company financial data. You are reviewing email messages from the various stakeholders for a project. The message from the security officer is shown in the Security Officer Email exhibit below.

TO: Database developer

From: Security Officer

Subject: SQL object requirements

We need to simplify the security settings for the SQL objects. Having a assign permissions at every object in SQL is tedious and leads to a problem. Documentation is also more difficult when we have to assign permissions at multiple levels. We need to assign the required permissions at one object, even though that object may be obtaining from other objects.

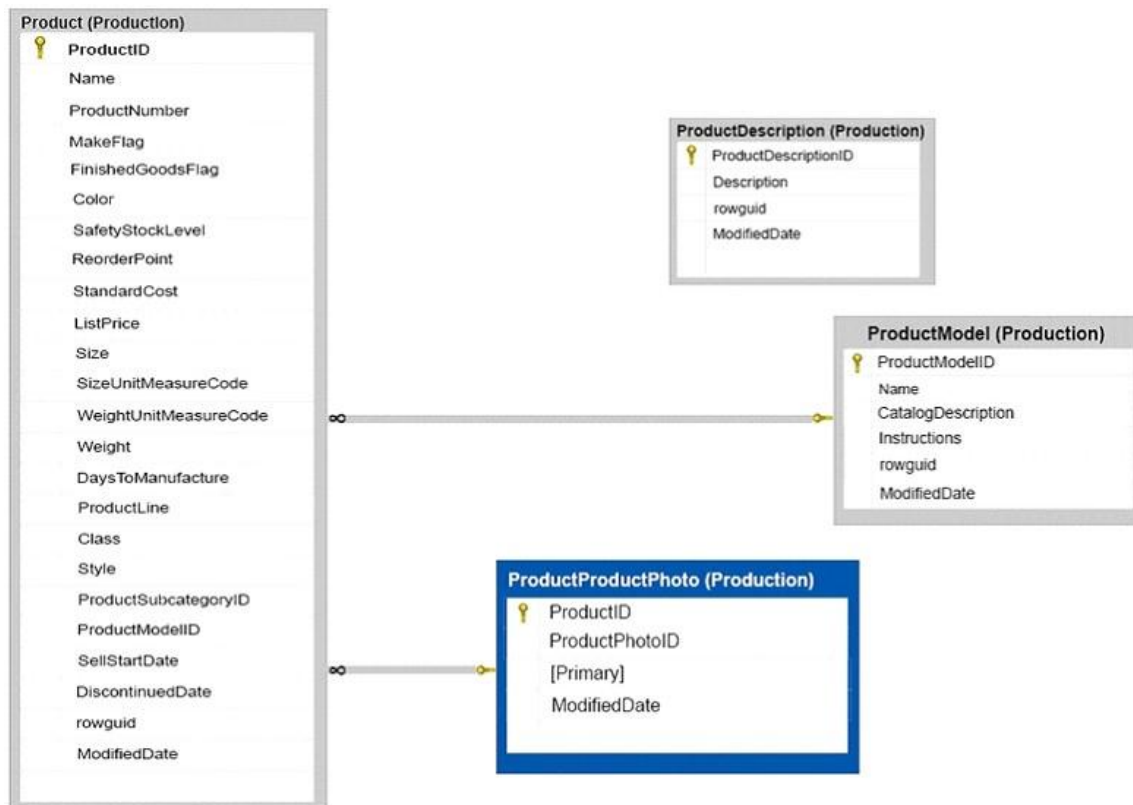
The message from the sales manager is shown in the Sales Manager Email exhibit below.

TO: Database developer

From: Sales Manager

Subject: Needed SQL objects

When creating objects for our use, they need to be flexible. We will be changing the base infrastructure frequently. We need components in SQL that will provide backward compatibility to our front end applications as the environments change so that do not need to modify the front end applications. We need objects that can provide a filtered set of the data. The data may be coming from multiple tables and we need an object that can provide access to all of the data through a single object reference. This is an example of the types of data we need to be able to have queries against without having to change the front end applications.



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

TO: Database developer

From: Web Developer

Subject: SQL Object component

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




Database Name: 

Table Name: 

Column Name: 

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

Answer Area

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

Answer:

Answer Area

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

QUESTION 17

You have a database that contains a table named Employees. The table stored information about the employees of your company. You need to implement the following auditing rules for the Employees table:

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

SQLution: You implement a user-defined function on the Employees table.

Does the SQLution meet the goal?

- A. Yes
- B. No

Answer: A

QUESTION 18

You have a database that contains a table named Employees. The table stored information about the employees of your company. You need to implement the following auditing rules for the Employees table:

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

SQLution: You implement a check constraint on the Employees table.

Does the SQLution meet the goal?

- A. Yes
- B. No

Answer: B

QUESTION 19

You have a database that contains a table named Employees. The table stored information about the employees of your company. You need to implement the following auditing rules for the Employees table:

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

SQLution: You implement a stored procedure on the Employees table.
Does the SQLution meet the goal?

- A. Yes
- B. No

Answer: B

QUESTION 20

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.

SQLution: 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 SQLution meet the goal?

- A. Yes
- B. No

Answer: B

QUESTION 21

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.

SQLution: You run the following Transact-SQL statement:

```
CREATE TABLE [dbo].[EmployeeAttendance] (  
    EmployeeID int NOT NULL,  
    DateChekedIn datetime2 NOT NULL,  
    DateCheclOut datetime2 NOT NULL)
```

Does the SQLution meet the goal?

- A. Yes
- B. No

Answer: B

QUESTION 22

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.

SQLution: You run the following Transact-SQL statement:

```
CREATE TABLE [dbo].[EmployeeAttendance] (  
    EmployeeID int NOT NULL,  
    DateChckedIn datetimeoffset(0) NOT NULL,  
    DateCheclOut datetimeoffset(0) NOT NULL)
```

Does the SQLution meet the goal?

- A. Yes
- B. No

Answer: A

QUESTION 23

The Account table was created 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.

SQLution: You run the following Transact-SQL statement:

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

Does the SQLution meet the goal?

- A. Yes
- B. No

Answer: B

QUESTION 24

The Account table was created 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.

SQLution: You run the following Transact-SQL statement:

```
CREATE CLUSTERED INDEX PK_Account ON Account(ProductCode);
```

Does the SQLution meet the goal?

- A. Yes
- B. No

Answer: B

QUESTION 25

The Account table was created 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.

SQLlution: 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 SQLlution meet the goal?


- A. Yes
- B. No

Answer: A

QUESTION 26

Drag and Drop Question

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. Changes to the price of any product must be less a 25 percent increase from the current price. The shipping department must be notified about order and shipping details when an order is entered

into the database. You need to implement the appropriate table objects. Which object should you use for each table? To answer, drag the appropriate objects to the correct tables. Each object 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.

Objects		Answer Area	
		Table	Objects
Foreign key constraint	Instead of trigger	Orders	
Check constraint	Primary key constraint	Products	
Unique constraint	After insert trigger		

Answer:

Answer Area	
Table	Objects
Orders	Foreign key constraint
Products	Primary key constraint

QUESTION 27

Hotspot Question

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 *	Products *
* OrderID	ProductID
ProductName	ProductName
ProductID	Description
EmployeeID	QtyonHand
OrderDate	SupplierName
	SupplierID
	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 implement a stored procedure that deletes a discontinued product from the Products table. You identify the following requirements. What should you do? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

Requirement

Transact-SQL segment

Handle errors

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

Display error message

```
ERROR MESSAGE()
PRINT
RAISERROR
RETURN
```

Answer:

Answer Area

Requirement

Transact-SQL segment

Handle errors

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

Display error message

```
ERROR MESSAGE()
PRINT
RAISERROR
RETURN
```

QUESTION 28

Hotspot Question

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 *	
* OrderID	
ProductName	
ProductID	
EmployeeID	
OrderDate	

Products *	
ProductID	
ProductName	
Description	
QtyonHand	
SupplierName	
SupplierID	
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. In the table below, identify the trigger types that meet the requirements. NOTE: Make only selection in each column. Each correct selection is worth one point.

Answer Area

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

Answer:

Answer Area

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

QUESTION 29

Hotspot Question

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 *	
*	OrderID
	ProductName
	ProductID
	EmployeeID
	OrderDate

Products *	
	ProductID
	ProductName
	Description
	QtyonHand
	SupplierName
	SupplierID
	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. The Leads table must include the columns described in the following table.

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

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

Answer Area

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

Answer:

Answer Area

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

QUESTION 30

Hotspot Question

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 *	Products *
* OrderID	ProductID
ProductName	ProductName
ProductID	Description
EmployeeID	QtyonHand
OrderDate	SupplierName
	SupplierID
	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 modify the database design to meet the following requirements. 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"/>

Answer:

Answer Area

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

QUESTION 31

Drag and Drop Question

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

Orders *	Products *
* OrderID	ProductID
ProductName	ProductName
ProductID	Description
EmployeeID	QtyonHand
OrderDate	SupplierName
	SupplierID
	Discontinued

The Customer table includes a column that stores the date for the last order that the customer placed. You plan to create a table named Leads. The Leads table is expected to contain approximately 20,000 records. Storage requirements for the Leads table must be minimized. You need to begin to modify the table design to adhere to third normal form. Which column should you remove for each table? To answer? drag the appropriate column names to the correct locations. Each column name may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Columns	Answer Area
ProductID	
ProductName	
Description	
EmployeeID	
OrderDate	
SupplierName	
SupplierID	
Discontinued	

Table	Column to remove
Products	Column
Orders	Column

Answer:

Answer Area

Table	Column to remove
Products	SupplierName
Orders	ProductName

QUESTION 32

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:

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

ProductReview table.

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:

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.

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 Sale Price column to values greater than four.

Finance department users must be able to retrieve data from the SalesHistory table for sales persons where the value of the SalesYTD column is above a certain threshold. You plan to create a memory-optimized table named SalesOrder. The table must meet the following requirements:

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

Performance for queries against the SalesOrder table that use where clauses with exact equality operations must be optimized. You need to enable referential integrity for the ProductReview table. How should you complete the relevant Transact-SQL statement? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

```
Alter Table dbo.ProductReview
WITH CHECK
ADD CONSTRAINT FK_productReview_Product FOREIGN KEY (ProductID)
REFERENCES Product (productID)
ON DELETE NO ACTION ON UPDATE CASCADE
```

Select two alternatives.

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

Answer: BC

QUESTION 33

Hotspot Question

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:

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

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:

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.

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 Sale Price column to values greater than four.

Finance department users must be able to retrieve data from the SalesHistory table for sales persons where the value of the SalesYTD column is above a certain threshold. You plan to create a memory-optimized table named SalesOrder. The table must meet the following requirements:

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

Performance for queries against the SalesOrder table that use where clauses with exact equality operations must be optimized. You need to create an object that allows finance users to be able to retrieve the required data. The object must not have a negative performance impact. How should you complete the Transact-SQL statements? To answer, select the appropriate TransactSQL segments in the answer area.

```
CREATE FUNCTION Sales.YIDSalesByPerson
  (@SalesPersonID int)
  RETURNS TABLE
  WITH SCHEMABINDING
  WITH ENCRYPTION
  RETURNS INT
  FROM Sales.SalesPerson
  WHERE SalesYTD > @minYIDSales
  ORDER BY SalesYTD desc);
```

Answer:

```
CREATE FUNCTION Sales.YIDSalesByPerson
  (@SalesPersonID int)
  RETURNS TABLE
  WITH SCHEMABINDING
  WITH ENCRYPTION
  RETURNS INT
  FROM Sales.SalesPerson
  WHERE SalesYTD > @minYIDSales
  ORDER BY SalesYTD desc);
```

QUESTION 34

You have a reporting database that includes a non-partitioned fact table named Fact_Sales. The table is persisted on disk. Users report that their queries take a long time to complete. The system administrator reports that the table takes too much space in the database. You observe that there are no indexes defined on the table, and many columns have repeating values. You need to create the most efficient index on the table, minimize disk storage and improve 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 column store index on the table.
- E. Create a nonclustered column store index on the table.
- F. Create a hash index on the table.

Answer: D

QUESTION 35

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 column store index on the table.
- E. Create a nonclustered column store index on the table.
- F. Create a hash index on the table.

Answer: C

QUESTION 36

Drag and Drop Question

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 level may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Isolation levels

read committed

serializable

read uncommitted

repeatable read

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

Answer:

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

read committed

read uncommitted

serializable

QUESTION 37

Drag and Drop Question

You have two database tables. Table1 is a partitioned table and Table 2 is a nonpartitioned table. Users report that queries take a long time to complete. You monitor queries by using Microsoft SQL Server Profiler. You observe lock escalation for Table1 and Table 2. You need to allow escalation of Table1 locks to the partition level and prevent all lock escalation for Table2. Which Transact-SQL statement should you run for each table? To answer, drag the appropriate Transact-SQL statements to the correct tables. Each command may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Transact-SQL statements

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

Answer Area

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

Answer:

Table	Transaction
Table1	ALTER TABLE Table1 SET (LOCK_ESCALATION - AUTO)
Table2	ALTER TABLE Table2 SET (LOCK_ESCALATION - DISABLE)

QUESTION 38

Drag and Drop Question

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 SQLution? To answer, move the appropriate Transact-SQL segments to the answer area and arrange then in the correct order.

Transact-SQL segments

Answer Area

```
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:

Answer Area

```
SET IMPLICIT_TRANSACTIONS ON
```

```
BEGIN TRAN
```

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

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

QUESTION 39

Hotspot Question

You are profiling a frequently used database table named UserEvents. The READ_COMMITTED_SNAPSHOT database option is set to OFF. In the trace results, you observe that lock escalation occurred for one stored procedure even though the number of locks in the database did not exceed memory or configuration thresholds. Events details are provided in the following table.

Attribute	Value
EventClass	LockEscalation
Mode	5 - X (X represents exclusive)
ObjectID	274100017 (objectID of table UserEvents)
Type	5 - OBJECT
TextData	<pre>CREATE PROCEDURE uspDeleteEvents (@EventType tinyint) AS BEGIN SET NOCOUNT ON; SET TRAN ISOLATION LEVEL READ COMMITTED; LABEL_DELETE: DELETE FROM UserEvents WITH (ROWLOCK) WHERE EventType = @EventType; END;</pre>

You need to modify the uspDeleteEvents stored procedure to avoid lock escalation. How should you modify the stored procedure? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

```
CREATE PROCEDURE uspDeleteEvents (@EventType tinyint)
AS
BEGIN
    SET NOCOUNT ON;
    SET TRAN ISOLATION LEVEL READ COMMITTED;
    LABEL_DELETE_DELETE_TOP(4000) FROM UserEvents WITH (ROWLOCK) WHERE EventType = @EventType
    LABEL_DELETE_DELETE FROM UserEvents WITH (PAGLOCK) WHERE EventType = @EventType
    LABEL_DELETE_DELETE TOP(400000) FROM UserEvents WITH (ROWLOCK) WHERE EventType = @EventType
    IF @@TRANCOUNT > 0 COMMIT TRAN
    WHILE @@ROWCOUNT > 0 GOTO LABEL_DELETE
    IF @@TRANCOUNT > 4000 COMMIT TRAN
    SET ROWCOUNT 4000
END;
```

Answer:

Answer Area

```
CREATE PROCEDURE uspDeleteEvents (@EventType tinyint)
AS
BEGIN
    SET NOCOUNT ON;
    SET TRAN ISOLATION LEVEL READ COMMITTED;
    LABEL_DELETE_DELETE_TOP(4000) FROM UserEvents WITH (ROWLOCK) WHERE EventType = @EventType
    LABEL_DELETE_DELETE FROM UserEvents WITH (PAGLOCK) WHERE EventType = @EventType
    LABEL_DELETE_DELETE TOP(400000) FROM UserEvents WITH (ROWLOCK) WHERE EventType = @EventType
    IF @@TRANCOUNT > 0 COMMIT TRAN
    WHILE @@ROWCOUNT > 0 GOTO LABEL_DELETE
    IF @@TRANCOUNT > 4000 COMMIT TRAN
    SET ROWCOUNT 4000
END;
```

QUESTION 40

You have a database that is experiencing deadlock issues when users run queries. You need to ensure that all deadlocks are recorded in XML format. What should you do?

- A. Create a Microsoft SQL Server Integration Services package that uses sys.dm_tran_locks.
- B. Enable trace flag 1224 by using the Database Consistency Checker(BDCC).
- C. Enable trace flag 1222 in the startup options for Microsoft SQL Server.
- D. Use the Microsoft SQL Server Profiler Lock: Deadlock event class.

Answer: C

QUESTION 41

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

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

The READ_COMMITTED_SNAPSHOT database option is set to OFF, and auto-content is set to ON. Within the stored procedures, no explicit transactions are defined. If JobB starts before JobA, it can finish in seconds. If JobA starts first, JobB takes a long time to complete. You need to use Microsoft SQL Server Profiler to determine whether the blocking that you observe in JobB is caused by locks acquired by JobA. Which trace event class in the Locks event category should you use?

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

Answer: A

QUESTION 42

Hotspot Question

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 ENCRYPTPRION 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.

Module	Object type
Module 1	<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>
Module 2	<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>

Answer:

Module	Object type
Module 1	<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>
Module 2	<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>

QUESTION 43

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

- At each level of the query plan, a low total number of rows are processed.
- The query uses many operations.

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

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

Answer: D

QUESTION 44

You are experiencing performance issues with the database server. You need to evaluate schema locking issues, plan cache memory pressure points, and backup I/O problems. What should you create?

- A. a System Monitor report
- B. a sys.dm_exec_query_stats dynamic management view query
- C. a sys.dm_exec_session_wait_stats dynamic management view query
- D. an Activity Monitor session in Microsoft SQL Management Studio

Answer: C

QUESTION 45

Hotspot Question

You are maintaining statistics for a database table named tblTransaction. The table contains more than 10 million records. You need to create a stored procedure that meets the following requirements:

- On weekdays, update statistics for a sample of the total number of records in the table.
- On weekends, update statistics by sampling all rows in the table.

A maintenance task will call this stored procedure daily. How should you complete the stored procedure? To answer, select the appropriate Transact-SQL segments in the answer area. NOTE: Each correct selection is worth one point.

```
CREATE PROCEDURE uspUpdateTxnStats
```

```
AS
```

```
BEGIN
```

```
    SET NOCOUNT ON;
```

```
    SET DATEFIRST 1;
```

```
    DECLARE @isWeekDay bit;
```

```
    SELECT @isWeekDay = CASE WHEN DATEPART (dw, GETDATE()) <=5 THEN 1 ELSE 0 END;
```

```
    IF @isWeekDay = 1
```

```
        BEGIN
```

UPDATE STATISTICS
SET STATISTICS
UPDATE TOP(20) STATISTICS
UPDATE #STATISTICS

tblTransaction

WITH FULLSCAN
WITH RESAMPLE
WITH SAMPLE 20 PERCENT
WITH SAMPLE 200000 ROWS
SELECT TOP(20) PERCENT

```
    END
```

```
ELSE
```

```
    BEGIN
```

UPDATE STATISTICS
SET STATISTICS
UPDATE TOP(20) STATISTICS
UPDATE #STATISTICS

tblTransaction

WITH FULLSCAN
WITH RESAMPLE
WITH SAMPLE 20 PERCENT
WITH SAMPLE 200000 ROWS
SELECT TOP(20) PERCENT

```
END;
```

```
END;
```


Answer:

```
CREATE PROCEDURE uspUpdateTxnStats
AS
BEGIN
    SET NOCOUNT ON;
    SET DATEFIRST 1;
    DECLARE @isWeekDay bit;
    SELECT @isWeekDay = CASE WHEN DATEPART (dw, GETDATE()) <=5 THEN 1 ELSE 0 END;
    IF @isWeekDay = 1
    BEGIN
```

UPDATE STATISTICS
SET STATISTICS
UPDATE TOP(20) STATISTICS
UPDATE #STATISTICS

tblTransaction
WITH FULLSCAN
WITH RESAMPLE
WITH SAMPLE 20 PERCENT
WITH SAMPLE 200000 ROWS
SELECT TOP(20) PERCENT

```
END
ELSE
BEGIN
```

UPDATE STATISTICS
SET STATISTICS
UPDATE TOP(20) STATISTICS
UPDATE #STATISTICS

tblTransaction
WITH FULLSCAN
WITH RESAMPLE
WITH SAMPLE 20 PERCENT
WITH SAMPLE 200000 ROWS
SELECT TOP(20) PERCENT

```
END;
END;
```

QUESTION 46

Drag and Drop Question

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 SQLution? 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 @duration Filter bigint = 100;
DECLARE @databaseId int;
SELECT @databaseId = DB_ID( 'MyDatabase');
```

Transact-SQL segments

Answer Area

```
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:**

```
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') ;
```

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

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

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

QUESTION 47

Hotspot Question

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 TransactSQL segments in the 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
         sys.dm_exec_requests
         sys.dm_db_index_usage_stats
         sys.dm_db_missing_index_details

INNER JOIN sys.indexes a ON b.object_id = a.object_id
         sys.tables
         sys.databases
         sys.views

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

Answer:

```
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
         sys.dm_exec_requests
         sys.dm_db_index_usage_stats
         sys.dm_db_missing_index_details

INNER JOIN sys.indexes a ON b.object_id = a.object_id
         sys.tables
         sys.databases
         sys.views

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

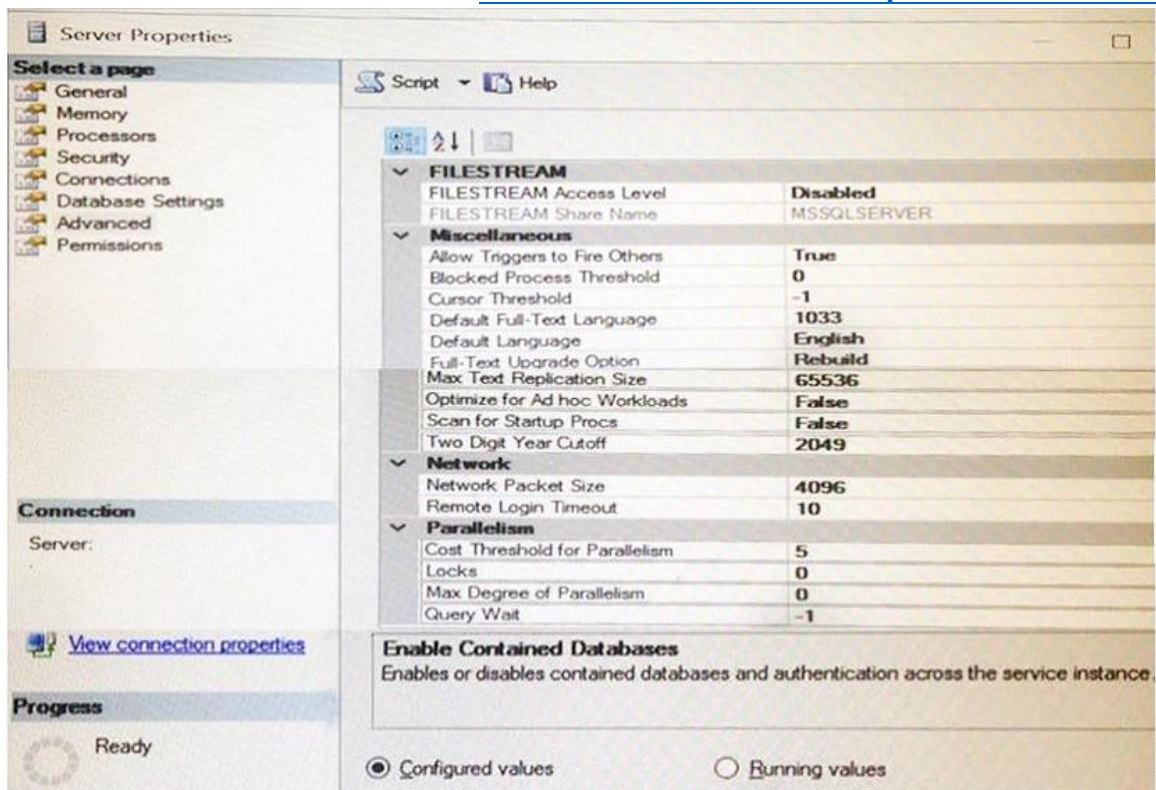
QUESTION 48

Hotspot Question

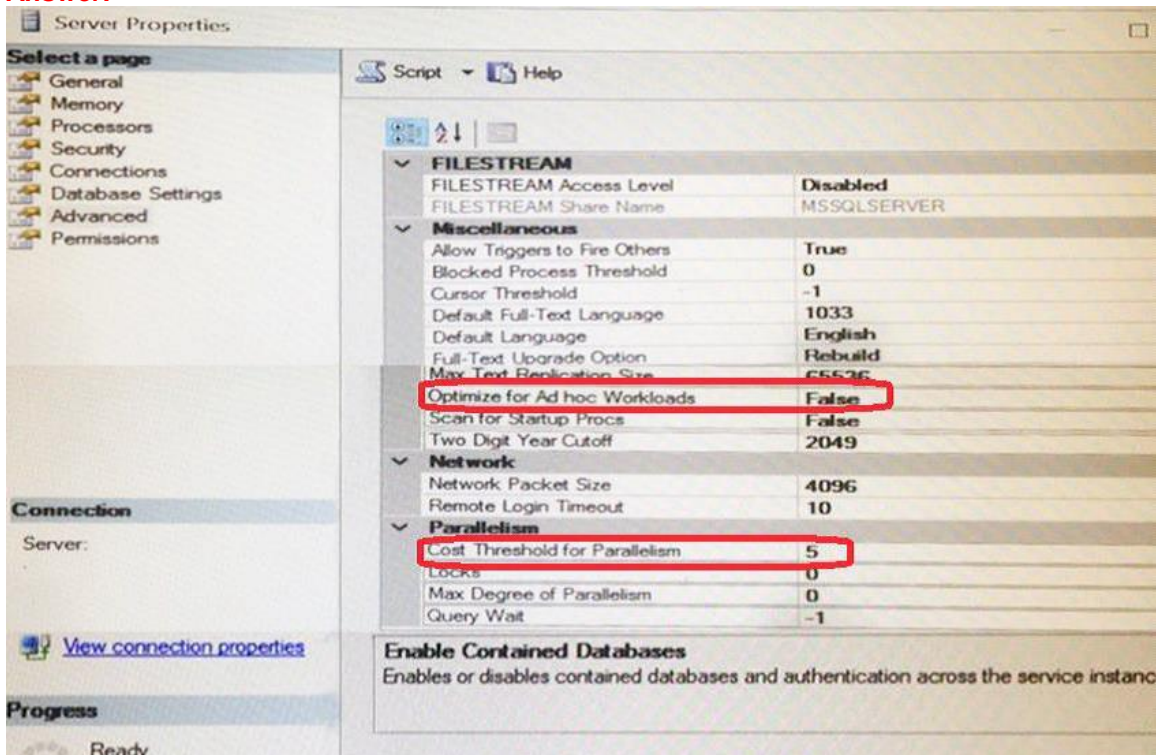
You are reviewing the execution plans in the query plan cache. You observe the following:

- There are a large number of single use plans.
- There are a large number of simple execution plans that use multiple CPU cores.

You need to configure the server to optimize query plan execution. Which two setting should you modify on the properties page for the Microsoft SQL Server instance? To answer, select the appropriate settings in the answer area.



Answer:



QUESTION 49

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 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.

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

QUESTION 50

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 inserted successfully into the Customers table.
- Return a value of 1 if data is not inserted successfully into the Customers table.
- Support TRY ... CATCH error handling.
- Be written by using Transact-SQL statements.

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: C

QUESTION 51

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

- Launch when table data is modified.
- Evaluate the state a table before and after a data modification and take action based on the difference.
- Prevent malicious or incorrect table data operations.
- Prevent changes that violate referential integrity by cancelling the attempted data modification.
- Run managed code packaged in an assembly that is created in the Microsoft.NET Framework and located into Microsoft SQL Server.

What should you create?

- A. extended procedure
- B. CLR procedure

- C. user-defined procedure
- D. DML trigger
- E. scalar-valued function
- F. table-valued function

Answer: D

QUESTION 52

You are developing an application to track customer sales. You need to return the sum of orders that have been finalized, given a specified order identifier. This value will be used in other Transact-SQL statements. You need to create a database object. 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: E

QUESTION 53

Hotspot Question

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:

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

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:

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.

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 Sale Price column to values greater than four.

Finance department users must be able to retrieve data from the SalesHistory table for sales persons where the value of the SalesYTD column is above a certain threshold. You plan to create a memory-optimized table named SalesOrder. The table must meet the following requirements:

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

Performance for queries against the SalesOrder table that use Where clauses with exact equality operations must be optimized. You need to create the Sales Order table. How should you complete the table definition? To answer? select the appropriate Transact-SQL segments in the answer area.

Answer Area

```
CREATE TABLE dbo.SalesOrder
(
    SalesOrderID int Primary KEY NONSELECTED HASH WITH (BUCKET_COUNT = 10000000),
    SalesOrderCode char(5) NOT NULL INDEX IX_SalesOrder
    EmployeeID int not null,
    CustomerID int not null,
    SalesAmount money
)
WITH
    (Memory_OPTIMIZED = ON),
    DURABILITY =
```

SCHEMA_ONLY
SCHEMA_AND_DATA

CLUSTERED
NONCLUSTERED
NONCLUSTERED HASH WITH (BUCKET_COUNT = 10000)
NONCLUSTERED HASH WITH (BUCKET_COUNT = 10000000)

Answer:

Answer Area

```
CREATE TABLE dbo.SalesOrder
(
    SalesOrderID int Primary KEY NONSELECTED HASH WITH (BUCKET_COUNT = 10000000),
    SalesOrderCode char(5) NOT NULL INDEX IX_SalesOrder
    EmployeeID int not null,
    CustomerID int not null,
    SalesAmount money
)
WITH
    (Memory_OPTIMIZED = ON),
    DURABILITY =
```

SCHEMA_ONLY
SCHEMA_AND_DATA

CLUSTERED
NONCLUSTERED
NONCLUSTERED HASH WITH (BUCKET_COUNT = 10000)
NONCLUSTERED HASH WITH (BUCKET_COUNT = 10000000)

QUESTION 54

Hotspot Question

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:

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

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:

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.

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 Sale Price column to values greater than four.

Finance department users must be able to retrieve data from the SalesHistory table for sales persons where the value of the SalesYTD column is above a certain threshold. You plan to create a memory-optimized table named SalesOrder. The table must meet the following requirements:

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

Performance for queries against the SalesOrder table that use Where clauses with exact equality operations must be optimized. You need to create a stored procedure named spDeleteCategory to delete records in the database. The stored procedure must meet the following requirements:

1. Delete records in both the BlogEntry and BlogCategory tables where CategoryId equals parameter @CategoryId.
2. Avoid locking the entire table when deleting records from the BlogCategory table.
3. If an error occurs during a delete operation on either table, all changes must be rolled back, otherwise all changes should be committed.

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

```
CREATE PROCEDURE spDeleteCategory
(@CategoryId int)
AS
BEGIN
    SET NOCOUNT ON;

    SET IMPLICIT_TRANSACTIONS ON
    SET IMPLICIT_TRANSACTIONS OFF
    SET TRANSACTION ISOLATION LEVEL READ COMMITTED
    SET TRANSACTION ISOLATION LEVEL SNAPSHOT

    BEGIN TRY
        DELETE FROM BlogEntry WHERE CategoryID = @CategoryId;
        ...
        DELETE FROM BlogCategory
        WITH (  ) WHERE CategoryId = @CategoryId;

        IF @@TRANCOUNT > 0  TRANSACTION;
        BE 
        BEGIN
        COMMIT

    END TRY
    BEGIN CATCH
        IF @@TRANCOUNT > 0  TRANSACTION;
        BEGIN
        COMMIT
        ROLLBACK

    END
```

Answer:

```
CREATE PROCEDURE spDeleteCategory  
(@CategoryID int)  
AS  
BEGIN
```

```
    SET NOCOUNT ON;
```

```
    SET IMPLICIT_TRANSACTIONS ON  
    SET IMPLICIT_TRANSACTIONS OFF  
    SET TRANSACTION ISOLATION LEVEL READ COMMITTED  
    SET TRANSACTION ISOLATION LEVEL SNAPSHOT
```

```
    BEGIN TRY  
        DELETE FROM BlogEntry WHERE CategoryID = @CategoryID;  
        ...  
        DELETE FROM BlogCategory  
        WITH (  ) WHERE CategoryId = @CategoryId;
```

ROWLOCK
 TABLOCKX

```
    IF @@TRANCOUNT > 0  
    BEGIN TRANSACTION;
```

BEGIN

COMMIT

```
    END TRY
```

```
    BEGIN CATCH
```

```
        IF @@TRANCOUNT > 0  
        BEGIN TRANSACTION;
```

BEGIN

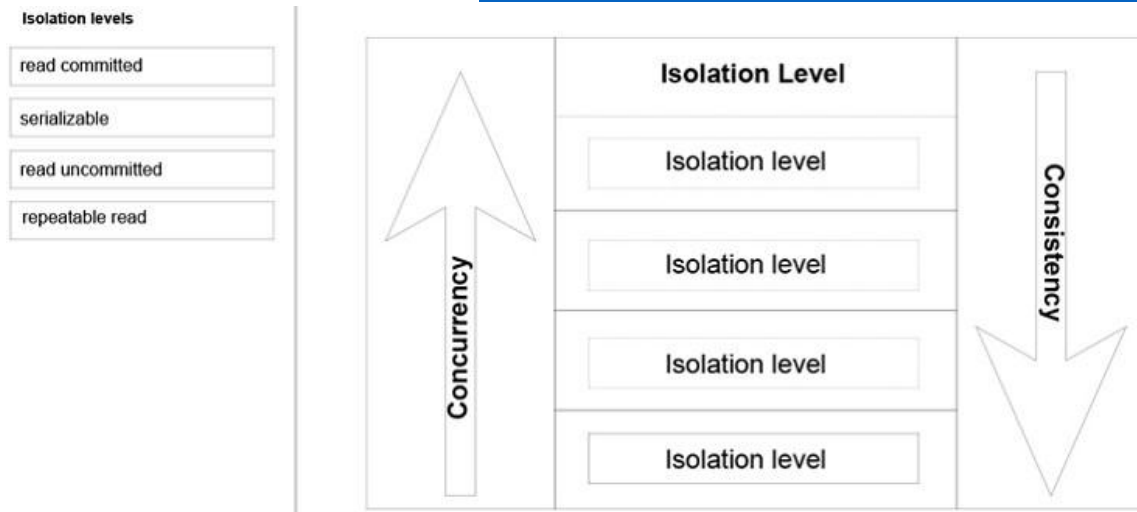
COMMIT

ROLLBACK

QUESTION 55

Drag and Drop Question

You are analyzing the performance of a database environment. Applications that access the database are experiencing locks that are held for a large amount of time. You are experiencing iSQLation phenomena such as dirty, nonrepeatable and phantom reads. You need to identify the impact of specific transaction iSQLation levels on the concurrency and consistency of data. What are the consistency and concurrency implications of each transaction iSQLation level? To answer, drag the appropriate iSQLation levels to the correct locations. Each iSQLation level may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.



Answer:



QUESTION 56

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


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

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

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

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

1. Allow users access to the CustomerName and CustomerNumber columns for active customers.
2. Allow changes to the columns that the view references. Modified data must be visible through the view.
3. Prevent the view from being published as part of Microsoft SQL Server replication.

Sales.FemaleCustomers must meet the following requirements:

1. Allow users access to the CustomerName, Address, City, State and PostalCode columns.
2. Prevent changes to the columns that the view references.
3. Only allow updates through the views that adhere to the view filter.

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

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

You run the spUpdateCustomerSummary stored procedure to make changes to customer account summaries. Other stored procedures call the spDeleteCustAcctRelationship to delete records from the CustomerToAccountBridge table. You must update the design of the Customer table to meet the following requirements.

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

You need to implement the design changes while minimizing data redundancy. What should you do?

- A. Split the table into three separate tables. Include the AccountNumber and CustomerID columns in the first table. Include the CustomerName and Gender columns in the second table. Include the AccountStatus column in the third table.
- B. Split the table into two separate tables. Include AccountNumber, CustomerID, CustomerName and Gender columns in the first table. Include the AccountNumber and AccountStatus columns in the second table.
- C. Split the table into two separate tables, Include the CustomerID and AccountNumber columns in the first table. Include the AccountNumber, AccountStatus, CustomerName and Gender columns in the second table.
- D. Split the table into two separate tables, Include the CustomerID, CustomerName and Gender columns in the first table. Include AccountNumber, AccountStatus and CustomerID columns in the second table.

Answer: D

QUESTION 57

.....

Get Complete Version Exam 70-762 Dumps with VCE and PDF Here



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