

70-762^{Q&As}

Developing SQL Databases

Pass Microsoft 70-762 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/70-762.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

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 create a file group for the indexes and a file group for the data files. You store the files for each file group on separate disks.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

Consider creating two additional File Groups: Tables and Indexes. It is best not to put your stuff in PRIMARY as that is where SQL SERVER stores all of its data and meta-data about your objects. You create your Table and Clustered Index (as that is the data for the table) on [Tables] and all Non-Clustered indexes on [Indexes].

QUESTION 2

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

You have a database named 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.

```
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)
);
```

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 Transacr-SQL statement:

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

You run the 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 the second update statement due to a CHECK constraint violation on the TotalDepositAccountCount column.

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

- A. 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.
- E. The value of the TotalDepositAccountCount column is increased.
- F. The statement that modifies TotalAccountCount column is excluded from the transaction.
- G. The value of the TotalDepositAccountCount column is decreased.

Correct Answer: D

QUESTION 3

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: Does the solution meet the goal?

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

A. Yes

B. No

Correct Answer: B

Datetime2 stores fractional seconds.

Datetime2 defines a date that is combined with a time of day that is based on a 24-hour clock. datetime2 can be considered as an extension of the existing datetime type that has a larger date range, a larger default fractional precision, and

optional user-specified precision.

References: <https://docs.microsoft.com/en-us/sql/t-sql/data-types/datetime2-transact-sql?view=sql-server-2017>

<https://msdn.microsoft.com/en-us/library/bb677335.aspx>

QUESTION 4

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

Correct Answer: A

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

QUESTION 5

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. a nonclustered index
- B. a schema-bound view
- C. a stored procedure
- D. an INSTEAD OF trigger

Correct Answer: B

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. Views or tables that participate in a view created with the SCHEMABINDING clause cannot be dropped unless that view is dropped or changed so that it no longer has schema binding.

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-view-transact-sql>

QUESTION 6

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 using the following Transact-SQL statement:

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

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

You frequently run the following Transact-SQL SELECT statements:

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

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>

QUESTION 7

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questions sets might have more than one correct solution,

while others might not have a correct solution.

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

You have a 3-TB database. The database server has 64 CPU cores.

You plan to migrate the database to Microsoft Azure SQL Database.

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

Solution: You select the Standard service tier.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

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>

QUESTION 8

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

Your company has employees in different regions around the world.

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

-Employee ID

-

date and time employee checked in to work

-

date and time employee checked out of work

Date and time information must be time zone aware and must not store fractional seconds.

Solution: You run the following Transact-SQL statement:


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

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Datetimeoffset, not datetimetype, 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>

QUESTION 9

You have a Microsoft Azure SQL Database. You enable Query Store for the database and configure the store to use the following settings:

SIZE_BASED_CLEANUP_MODE = OFF

STALE_QUERY_THRESHOLD_DAYS = 60

MAX_STORAGE_SIZE_MB = 100

QUERY_CAPTURE_MODE = ALL

You use Azure Query Performance Insight to review queries. You observe that new queries are not displayed after 15 days and that the Query Store is set to read-only mode.

If the Query Store runs low on data space, the store must prioritize queries that run regularly or queries that consume applicant resources.

You must set the Query Store to read_write mode and determine the performance of queries from the past 60 days.

Which three actions should you perform? Each correct step presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Set the value of the CLEANUP_POLICY setting to (STALE_QUERY_THRESHOLD_DAYS = 75)

B. Set the value of the QUERY_CAPTURE_MODE setting to AUTO

C. Increase the value for the MAX_STORAGE_SIZE_MB setting

D. Set the value of the SIZE_BASED_CLEANUP_MODE setting to AUTO

E. In the Azure portal, navigate to Query Performance Insight. Use the Custom tab to select a period of 2 months.

Correct Answer: BCD

B: Capture mode:

All ?Captures all queries. This is the default option.

Auto ?Infrequent queries and queries with insignificant cost are ignored. (Ad hoc recommended)

None ?Query Store stops capturing new queries.

C: Max Size (MB): Specifies the limit for the data space that Query Store can consume within the database. This is the most important setting that directly affects operation mode of the Query Store.

While Query Store collects queries, execution plans and statistics, its size in the database grows until this limit is reached. When that happens, Query Store automatically changes the operation mode to read-only and stops collecting new

data. You should monitor this closely to make sure you have sized the store appropriately to contain the full history you'd like to retain.

D: Size Based Cleanup Mode: Specifies whether automatic data cleanup will take place when Query Store data size approaches the limit.

It is strongly recommended to activate size-based cleanup to makes sure that Query Store always runs in read-write mode and collects the latest data.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store>

QUESTION 10

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

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

- Record any changes that are made to the data in the Employees table.

-

Customize the data recorded by the audit operations.

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

Does the solution meet the goal?

A.

Yes

B.

No

Correct Answer: A

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

made to user tables in a database.

Change data is made available to change data capture consumers through table-valued functions (TVFs).

References:<https://msdn.microsoft.com/en-us/library/cc645858.aspx>

QUESTION 11

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

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

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

What should you do?

- A. Use the NOLOCK option.
- B. Execute the 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.

Correct Answer: E

Snapshot isolation enhances concurrency for OLTP applications.

Once snapshot isolation is enabled, updated row versions for each transaction are maintained in tempdb. A unique transaction sequence number identifies each transaction, and these unique numbers are recorded for each row version. The transaction works with the most recent row versions having a sequence number before the sequence number of the transaction. Newer row versions created after the transaction has begun are ignored by the transaction. References: <https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/snapshot-isolation-in-sql-server>

QUESTION 12

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

Correct Answer: B

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>

QUESTION 13

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

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

- Record any changes that are made to the data in the Employees table.

-

Customize the data recorded by the audit operations.

Solution: You implement a stored procedure on the Employees table.

Does the solution meet the goal?

A.

Yes

B.

No

Correct Answer: B

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

QUESTION 14

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

Correct Answer: A

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>

QUESTION 15

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

Information and details provided in a question apply only to that question.

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

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

What should you do?

- A. Create `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.

Correct Answer: B

`sys.dm_exec_sessions` returns one row per authenticated session on SQL Server. `sys.dm_exec_sessions` is a server-scope view that shows information about all active user connections and internal tasks. This information includes client version, client program name, client login time, login user, current session setting, and more. Use `sys.dm_exec_sessions` to first view the current system load and to identify a session of interest, and then learn more information about that session by using other dynamic management views or dynamic management functions.

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

[70-762 PDF Dumps](#)

[70-762 VCE Dumps](#)

[70-762 Practice Test](#)