



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.pass4itsure.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**

You run the following Transact-SQL statement:

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

There are multiple unique OrderID values. Most of the UnitPrice values for the same OrderID are different.

You need to create a single index seek query that does not use the following operators: Nested loop Sort Key lookup

Which Transact-SQL statement should you run?

- A. CREATE INDEX IX_OrderLines_1 ON OrderLines (OrderID, UnitPrice) INCLUDE (Description, Quantity)
- B. CREATE INDEX IX_OrderLines_1 ON OrderLines (OrderID, UnitPrice) INCLUDE (Quantity)
- C. CREATE INDEX IX_OrderLines_1 ON OrderLines (OrderID, UnitPrice, Quantity)
- D. CREATE INDEX IX_OrderLines_1 ON OrderLines (UnitPrice, OrderID) INCLUDE (Description, Quantity)

Correct Answer: A

An index with nonkey columns can significantly improve query performance when all columns in the query are included in the index either as key or nonkey columns. Performance gains are achieved because the query optimizer can locate all

the column values within the index; table or clustered index data is not accessed resulting in fewer disk I/O operations.

Note: All data types except text, ntext, and image can be used as nonkey columns.

Incorrect Answers:

C: Redesign nonclustered indexes with a large index key size so that only columns used for searching and lookups are key columns.

D: The most unique column should be the first in the index.

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-index-transact-sql?view=sql-server-2017>

QUESTION 2

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 Performance statistics 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.error_reported event added.

Correct Answer: A

The Performance Statistics event class can be used to monitor the performance of queries, stored procedures, and triggers that are executing. Each of the six event subclasses indicates an event in the lifetime of queries, stored procedures, and triggers within the system. Using the combination of these event subclasses and the associated `sys.dm_exec_query_stats`, `sys.dm_exec_procedure_stats` and `sys.dm_exec_trigger_stats` dynamic management views, you can reconstitute the performance history of any given query, stored procedure, or trigger.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/event-classes/performance-statistics-event-class?view=sql-server-2017>

QUESTION 3

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

You have a database named dbReporting. Users run a large number of read-only ad hoc queries against the database. The application and all ad hoc queries use default database transaction isolation levels. You set the value of the

READ_COMMITTED_SNAPSHOT database option to ON.

You have an application that updates 10 tables sequentially and modifies a large volume of records in a single transaction. The updates are isolated from each other.

Users report an error which indicates that the version store is full.

You need to reduce the number of occurrences of the error.

Solution: You increase the maximum database size for the dbReporting database.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B



QUESTION 4

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 make 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 5

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

You are developing a new application that uses a stored procedure. The stored procedure inserts thousands of records as a single batch into the Employees table.

Users report that the application response time has worsened since the stored procedure was updated. You examine disk-related performance counters for the Microsoft SQL Server instance and observe several high values that include a disk performance issue. You examine wait statistics and observe an unusually high WRITELOG value.

You need to improve the application response time.

Solution: You add a unique clustered index to the Employees table.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

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

QUESTION 6

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 need to collect data from the following two sources:

1.

The performance counters of the operating system

2.



Microsoft SQL Server events

You must analyze the two datasets side-by side by using a single tool.

Solution: You use SQL Server Profiler and SQL Server Extended Events to collect performance data. You use SQL Server Profiler to analyze the data.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Reference: <https://www.sqlshack.com/dba-guide-sql-server-performance-troubleshooting-part-2-monitoring-utilities/>

QUESTION 7

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

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: A

Create a clustered index on theAccountNumber column as it is unique. Create a nonclustered index that includes the ProductCode column.

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

QUESTION 8

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

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

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

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



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

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

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.Female.Customers 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 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 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.

Correct Answer: D

Two tables is enough. CustomerID must be in both tables.

QUESTION 9

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

Your company has employees in different regions around the world. You need to create a database table that stores the following employee attendance information: Employee ID date and time employee checked in to work date and time employee checked out of work

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

Solution: You run the following Transact-SQL statement:

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



Does the solution meet the goal?

A. Yes

B. No

Correct Answer: A

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

Syntax: datetimeoffset [(fractional seconds precision)]

For the use "datetimeoffset(0)", the Fractional seconds precision is 0, which is required here.

References: <https://msdn.microsoft.com/en-us/library/bb630289.aspx>

QUESTION 10

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 need to collect data from the following two sources:

1.

The performance counters of the operating system

2.

Microsoft SQL Server events

You must analyze the two datasets side-by-side by using a single tool.

Solution: You use dynamic management views and Data Collector Sets (DCs) in Performance Monitor to collect performance data. You use SQL Server Management Studio (SSMS) to analyze the data.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Reference: <https://www.sqlshack.com/dba-guide-sql-server-performance-troubleshooting-part-2-monitoring-utilities/>

QUESTION 11

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.

QUESTION 12

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_tran_database_transaction` dynamic management view query
- C. an Extended Events session that uses Query Editor
- D. a Microsoft SQL Profiler trace

Correct Answer: C

Extended Events: considered as "the best way" by the SQL Server purists. You can configure Extended Events to find



Locking Issues in SQL Server. Incorrect Answers:

D: SQL Trace and SQL Server Profiler are deprecated. The Microsoft.SqlServer.Management.Trace namespace that contains the Microsoft SQL Server Trace and Replay objects are also deprecated.

This feature is in maintenance mode and may be removed in a future version of Microsoft SQL Server. Avoid using this feature in new development work, and plan to modify applications that currently use this feature.

Use Extended Events instead.

References:

<https://www.mssqltips.com/sqlservertip/5752/configuring-extended-events-to-find-locking-issues-in-sql-server/>

QUESTION 13

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

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

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

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

You need to examine delays in executed threads, including errors with queries and batches.

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.

Correct Answer: G

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

A: sys.dm_os_waiting_tasks returns information about the wait queue of tasks that are waiting on some resource.



References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-os-wait-stats-transact-sql>

QUESTION 14

You have a data warehouse fact table that has a clustered columnstore index.

You have multiple CSV files that contain a total of 3 million rows of data.

You need to upload the data to the fact table. The solution must avoid the delta group when you import the data.

Which solution will achieve the goal in the least amount of time?

- A. Load the source data to a staging table. Load the data to the fact table by using the INSERT_SELECT statement and specify the TABLOCK option on the staging table
- B. Create a Microsoft SQL Server Integration Services (SSIS) package. Use multiple data flow tasks to load the data in parallel.
- C. Load the source data to the fact table by running bcp.exe and specify the ?TABLOCK option
- D. Load the source data to the fact table by using the BULK INSERT statement and specify the TABLOCK option

Correct Answer: A

If you are loading data only to stage it before running more transformations, loading the table to heap table will be much faster than loading the data to a clustered columnstore table. In addition, loading data to a [temporary table][Temporary]

will also load much faster than loading a table to permanent storage.

A common pattern for data load is to load the data into a staging table, do some transformation and then load it into the target table using the following command

```
INSERT INTO
```

```
SELECT FROM
```

This command loads the data into the columnstore index in similar ways to BCP or Bulk Insert but in a single batch. If the number of rows in the staging table

loaded into compressed rowgroup. One key limitation was that this INSERT operation was single threaded. To load data in parallel, you could create multiple staging table or issue INSERT/SELECT with non-overlapping ranges of rows from

the staging table. This limitation goes away with SQL Server 2016 (13.x). The command below loads the data from staging table in parallel but you will need to specify TABLOCK.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-data-loading-guidance?view=sql-server-2017#plan-bulk-load-sizes-to-minimize-delta-rowgroups>

QUESTION 15



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>

[Latest 70-762 Dumps](#)

[70-762 VCE Dumps](#)

[70-762 Exam Questions](#)