

TestValid

LOGIN / REGISTER | MY CART (5) Q



HOME

CERTIFICATIONS

ABOUT

HOW TO PAY?

GUARANTEE

FAQ

WELCOME TO TESTVALID

Valid test questions and answers for IT certification actual test



With the help of our website, you will pass actual test by practicing test questions and remembering the test answers skillfully.



Quality and Value

TestValid Practice Exams are written to the highest standards of technical accuracy, using only certified subject matter experts and published authors for development - no all dumps.



Tested and Approved

We are committed to the process of vendor and third party approvals. We believe professionals and executives alike deserve the confidence of quality coverage these authorizations provide.



Easy to Pass

If you prepare for the exams using our TestValid testing engine, it is easy to succeed for all certifications in the first attempt. You don't have to deal with all dumps or any free torrent / rapidshare all stuff.



Try Before Buy

TestValid offers free demo of each product. You can check out the interface, question quality and usability of our practice exams before you decide to buy.

<http://www.testvalid.com>

Valid test questions and answers for IT certification actual test

Exam : **70-762**

Title : **Developing SQL Databases**

Vendor : **Microsoft**

Version : **DEMO**

NO.1 Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

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

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

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

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

There are no foreign key relationships between TBL1 and TBL2.

You need to minimize the amount of time required for the two queries to return records from the tables.

What should you do?

- A.** Create unique constraints on both TBL1 and TBL2. Create a partitioned view that combines columns from TBL1 and TBL2.
- B.** Create a nonclustered index on tbl2 only.
- C.** Create an indexed view that combines columns from TBL1 and TBL2.
- D.** Drop existing indexes on TBL1 and then create a clustered columnstore index. Create a nonclustered columnstore index on TBL1. Create a nonclustered index on TBL2.
- E.** Create check constraints on both TBL1 and tbl2. Create a partitioned view that combines columns from TBL1 and tbl2.
- F.** Drop existing indexes on TBL1 and then create a clustered columnstore index. Create a nonclustered columnstore index on TBL1. Make no changes to TBL2.
- G.** Create a clustered index on TBL1. Create a nonclustered index on tbl2 and add the most frequently queried columns as included columns.
- H.** Create clustered indexes on TBL1 and TBL2.

Answer: F

NO.2 You manage a database with tables named Invoice and InvoiceDetails. Each invoice may have multiple records.

Users update the InvoiceDetails table by using a .NET web application. The application retrieves records from both tables and updates the tables by running an inline update statement.

Users experience slow performance when updating records in the application. The solution must meet the following requirements:

- * Must use a stored procedure.
- * Must not use inline update statements
- * Must use a table-valued parameter.
- * Must call the stored procedure to update all records.

You need to optimize performance.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

- Create a stored procedure and use invDetails as a read-only input parameter.
- Create a stored procedure and use invDetails as the input parameter.
- Send invoice detail records to the stored procedure as a string value for the input parameter.
- Create a user-defined table type named invDetails like schema of InvoiceDetails table.
- Create an alias type from VARCHAR (MAX) Named invDetails.
- Send invoice detail records to the stored procedure as a ADO.NET table for the input parameter

Answer Area



Answer:

Actions

Create a stored procedure and use invDetails as a read-only input parameter.

Create a stored procedure and use invDetails as the input parameter.

Send invoice detail records to the stored procedure as a string value for the input parameter.

Create a user-defined table type named invDetails like schema of InvoiceDetails table.

Create an alias type from VARCHAR (MAX) Named invDetails.

Send invoice detail records to the stored procedure as a ADO.NET table for the input parameter

Answer Area

Create a user-defined table type named invDetails like schema of InvoiceDetails table.

Create a stored procedure and use invDetails as a read-only input parameter.

Send invoice detail records to the stored procedure as a ADO.NET table for the input parameter

Explanation

Answer Area

Create a user-defined table type named invDetails like schema of InvoiceDetails table.

Create a stored procedure and use invDetails as a read-only input parameter.

Send invoice detail records to the stored procedure as a ADO.NET table for the input parameter

Box 1: Create a user-defined table type...

Table-valued parameters are declared by using user-defined table types. You can use table-valued parameters to send multiple rows of data to a Transact-SQL statement or a routine, such as a stored procedure or function, without creating a temporary table or many parameters.

Box 2: ...read-only input parameter.

Table-valued parameters must be passed as input READONLY parameters to Transact-SQL routines.

Box 3:

Example

The following example uses Transact-SQL and shows you how to create a table-valued parameter

type, declare a variable to reference it, fill the parameter list, and then pass the values to a stored procedure.

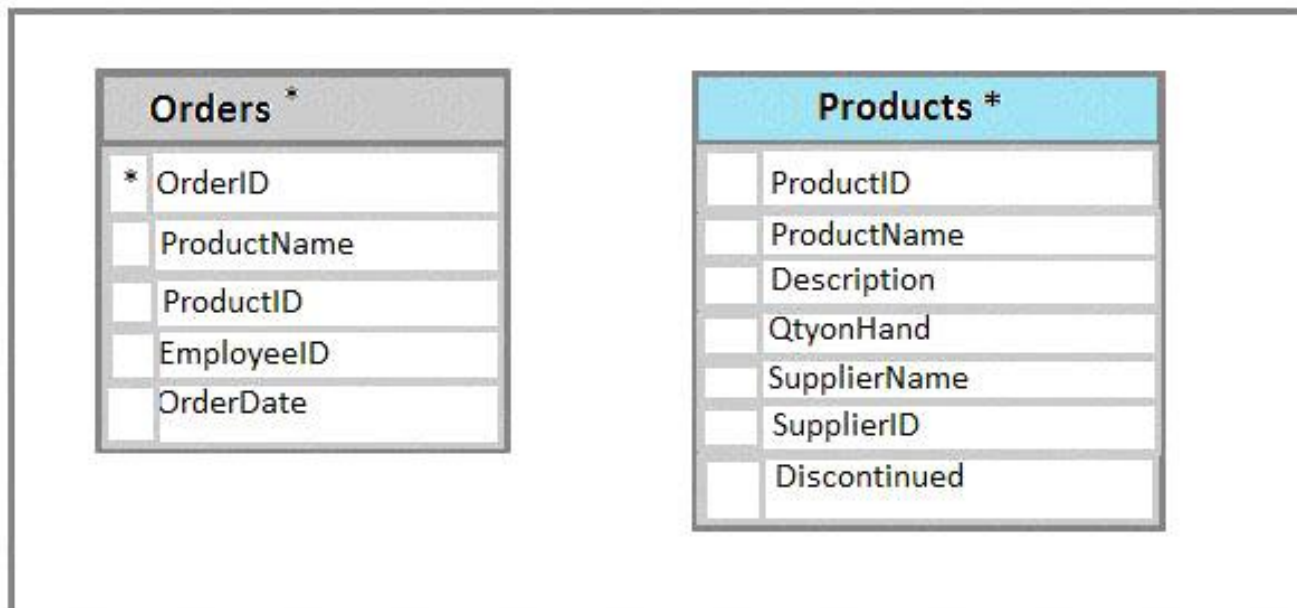
```
USE AdventureWorks2012;
/* Create a table type. */
CREATE TYPE LocationTableType AS TABLE
( LocationName VARCHAR(50)
, CostRate INT );
GO
/* Create a procedure to receive data for the table-valued parameter. */ CREATE PROCEDURE dbo.
usp_InsertProductionLocation
@TVP LocationTableType READONLY
Etc.
/* Declare a variable that references the type. */
DECLARE @LocationTVP AS LocationTableType;
/* Add data to the table variable. */
INSERT INTO @LocationTVP (LocationName, CostRate)
SELECT Name, 0.00
FROM AdventureWorks2012.Person.StateProvince;
/* Pass the table variable data to a stored procedure. */
EXEC usp_InsertProductionLocation @LocationTVP;
GO
```

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/tables/use-table-valued-parameters-database-engine?vie>

NO.3 Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in the series. 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.



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

- ProductID
- ProductName
- Description
- EmployeeID
- OrderDate
- SupplierName
- SupplierID
- Discontinued

Answer Area

Table	Column to remove
Products	Column
Orders	Column

Answer:

Columns

- ProductID
- ProductName
- Description
- EmployeeID
- OrderDate
- SupplierName
- SupplierID
- Discontinued

Answer Area

Table	Column to remove
Products	SupplierName
Orders	ProductName

Explanation

Answer Area

Table	Column to remove
Products	SupplierName
Orders	ProductName

In the Products table the SupplierName is dependent on the SupplierID, not on the ProductID.
In the Orders table the ProductName is dependent on the ProductID, not on the OrderID.

Note:

A table is in third normal form when the following conditions are met:

- * It is in second normal form.
- * All non-primary fields are dependent on the primary key.

Second normal form states that it should meet all the rules for First 1Normal Form and there must be no partial dependencies of any of the columns on the primary key.

First normal form (1NF) sets the very basic rules for an organized database:

- * Define the data items required, because they become the columns in a table. Place related data items in a table.
- * Ensure that there are no repeating groups of data.
- * Ensure that there is a primary key.

References: <https://www.tutorialspoint.com/sql/third-normal-form.htm>

NO.4 You create tables by using the following Transact-SQL statements:

```
CREATE TABLE Customer (
  CustomerId INT IDENTITY (1,1) PRIMARY KEY,
  Name VARCHAR (200) NOT NULL,
  CreditLimit DECIMAL (16, 2) NOT NULL
)
```

```
CREATE TABLE Address(
  AddressId INT IDENTITY (1, 1) PRIMARY KEY,
  CustomerId INT NOT NULL REFERENCES Customer (CustomerId),
  Address1 VARCHAR (200) NOT NULL,
  Address2 VARCHAR (200) NULL,
  City VARCHAR (20) NOT NULL,
  CountryCode CHAR (2) NOT NULL,
  IsActive BIT NOT NULL,
)
```

Each customer may have multiple addresses but only one is the primary address.

You must plan a solution that meets the following requirements:

- * Return both customers and address information.
- * Return only the primary address of the customer.
- * Allow updates of all customer information and address details with the exception of the identity columns and the IsActive column.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Create a view joining both tables and showing all columns with WHERE and WITH CHECK OPTION statements.	
Allow the application to query and update records directly from the table-valued function.	⊙
Create a table-valued function that returns all columns by joining both tables.	
Allow the application to query and update records directly from the view.	⊙
Create a view joining both tables and showing all columns with a WHERE statement.	
Add the INSTEAD OF UPDATE statement to the view.	
Create a view joining both tables and showing all columns.	

Answer:

Actions

- Create a view joining both tables and showing all columns with WHERE and WITH CHECK OPTION statements.
- Allow the application to query and update records directly from the table-values function.
- Create a table-valued function that returns all columns by joining both tables.
- Allow the application to query and update records directly from the view.
- Create a view joining both tables and showing all columns with a WHERE statement.
- Add the INSTEAD OF UPDATE statement to the view.
- Create a view joining both tables and showing all columns.

Answer Area

Create a view joining both tables and showing all columns with WHERE and WITH CHECK OPTION statements.

Add the INSTEAD OF UPDATE statement to the view.

Allow the application to query and update records directly from the table-values function.

Create a view joining both tables and showing all columns with WHERE and WITH CHECK OPTION statements.

Add the INSTEAD OF UPDATE statement to the view.

Allow the application to query and update records directly from the table-values function.

NO.5 You have a reporting application that uses a table named Table1. You deploy a new batch update process to perform updates to Table1.

The environment is configured with the following properties:

- * The database is configured with the default isolation setting.
- * The application and process use the default transaction handling.

You observe the application cannot access any rows that are in use by the process.

You have the following requirements:

- * Ensure the application is not blocked by the process.
- * Ensure the application has a consistent view of the data
- * Ensure the application does not read dirty data.

You need to resolve the issue and meet the requirements with the least amount of administrative effort.

What should you do?

- A.** Enable the database for the READ_COMMITTED_SNAPSHOT isolation level.
- B.** Enable the database for the ALLOW_SNAPSHOT_ISOLATION isolation level. Modify the application and the update process for the SNAPSHOT isolation level.
- C.** Enable the database for the ALLOW_SNAPSHOT_ISOLATION isolation level. Modify the application for the SERIALIZABLE isolation level.
- D.** Enable the application for the WITH (NOLOCK) hint.

Answer: A

Explanation

Snapshot isolation must be enabled by setting the ALLOW_SNAPSHOT_ISOLATION ON database option before it is used in transactions. This activates the mechanism for storing row versions in the temporary database (tempdb).

READ COMMITTED is the default isolation level for SQL Server. It prevents dirty reads by specifying that statements cannot read data values that have been modified but not yet committed by other transactions. Other transactions can still modify, insert, or delete data between executions of individual statements within the current transaction, resulting in non-repeatable reads, or "phantom" data.

NO.6 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.

The screenshot shows the 'Server Properties' dialog box for a Microsoft SQL Server instance. The 'Parallelism' section is expanded, showing the following settings:

Setting	Value
Cost Threshold for Parallelism	5
Locks	0
Max Degree of Parallelism	0
Query Wait	-1

The 'Miscellaneous' section is also expanded, showing the following settings:

Setting	Value
Allow Triggers to Fire Others	True
Blocked Process Threshold	0
Cursor Threshold	-1
Default Full-Text Language	1033
Default Language	English
Full-Text Upgrade Option	Rebuild
Max Text Replication Size	65536
Optimize for Ad hoc Workloads	False
Scan for Startup Procs	False
Two Digit Year Cutoff	2049

The 'Network' section shows the following settings:

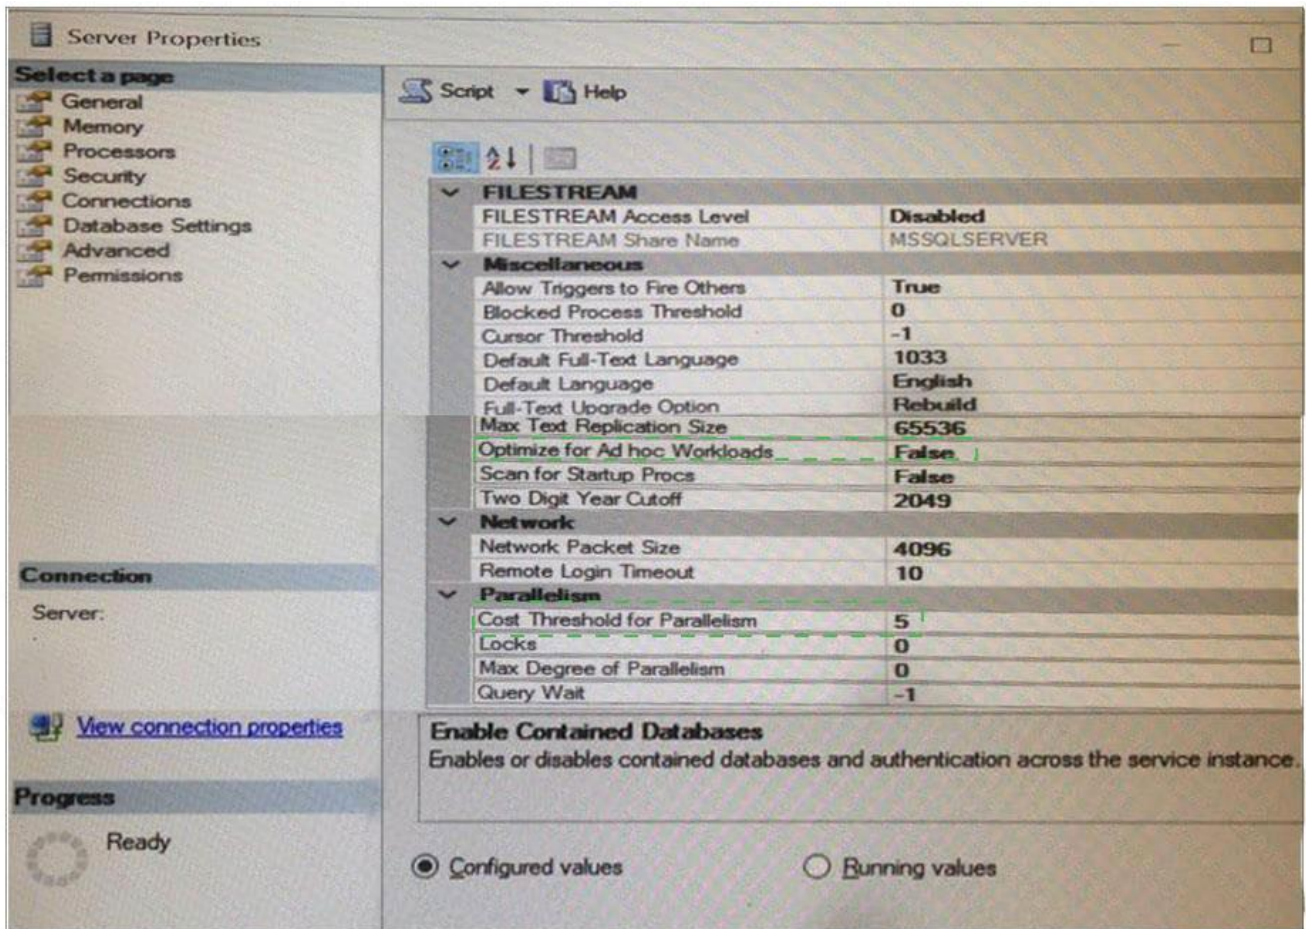
Setting	Value
Network Packet Size	4096
Remote Login Timeout	10

The 'FILESTREAM' section shows the following settings:

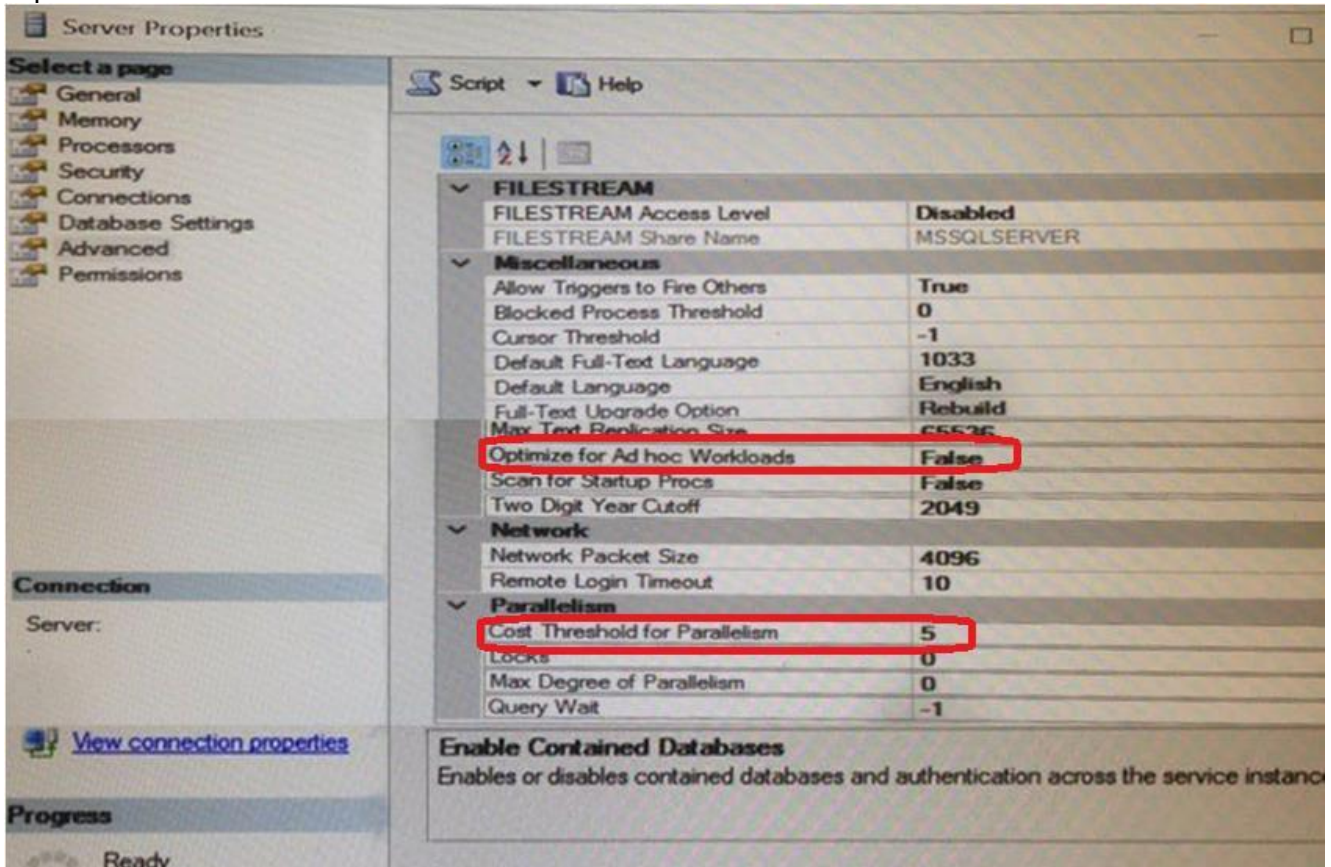
Setting	Value
FILESTREAM Access Level	Disabled
FILESTREAM Share Name	MSSQLSERVER

The 'Enable Contained Databases' section is at the bottom of the dialog, with the description: 'Enables or disables contained databases and authentication across the service instance.' At the bottom of the dialog, there are two radio buttons: 'Configured values' (selected) and 'Running values'.

Answer:



Explanation



* Optimize for ad hoc workloads

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

* Cost Threshold for Parallelism

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

5 means 5 seconds, but is 5 seconds on a machine internal to Microsoft from some time in the 1990s.

There's no way to relate it to execution time on your current machine, so we treat it as a pure number now.

Raising it to 50 is a common suggestion nowadays, so that more of your simpler queries run on a single thread.

NO.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. Determine whether the solution meets the stated goals.

You have a table that has a clustered index and a nonclustered index. The indexes use different columns from the table. You have a query named Users report that Query1 takes a long time to report results. You run and review the following statistics for an index seek operation:

Index Seek (NonClustered)

Scan a particular range of rows from a nonclustered index.

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

You need to resolve the performance issue.

Solution: You rebuild the clustered index.

Does the solution meet the goal?

A. No

B. Yes

Answer: A

Explanation

The query uses the nonclustered index, so improving the clustered index will not help.

We should update statistics for the nonclustered index.

NO.8 Note: This question is part of a series of questions that present the same scenario. Each question in this series contains a unique solution. Determine whether the solution meets the stated goals.

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

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

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

You frequently run the following Transact-SQL SELECT statements:

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

You must avoid table scans when you run the queries.

You need to create one or more indexes for the table.

Solution: You run the following Transact-SQL statement:

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

Does the solution meet the goal?

A. No

B. Yes

Answer: A

Explanation

We need an index on the productCode column as well.

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

NO.9 Database users report that SELECT statements take a long time to return results. You run the following Transact-SQL statement:

```
SELECT OBJECT_NAME([object_id]) AS [object_name],
d.equality_columns, d.inequality_columns, d.included_columns
FROM sys.dm_db_missing_index_details;
```

Object_name	Equality_columns	Inequality_columns	Included_columns
[Users]	[CountryCode]	[UserStatus]	[UserName]

You need to create one nonclustered covering index that contains all of the columns in the above table. You must minimize index key size.

Which Transact-SQL statement should you run?

- A.** CREATE NONCLUSTERED INDEX IX_User ON Users (UserStatus, CountryCode) INCLUDE (UserName);
- B.** CREATE NONCLUSTERED INDEX IX_User ON Users (CountryCode, UserStatus, UserName);
- C.** CREATE NONCLUSTERED INDEX IX_User ON Users (CountryCode, UserStatus) INCLUDE (UserName);
- D.** CREATE NONCLUSTERED INDEX IX_User ON Users (CountryCode, UserName);

Answer: C

Explanation

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/create-indexes-with-included-columns>

NO.10 Note: This question is part of a series of questions that use the same or similar answer choices. An Answer choice may be correct for more than one question in the series. Each question independent of the other questions in this series. Information and details provided in a question apply only to that question.

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

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

You need to create a baseline set of metrics to report how the computer running SQL Server operates under normal load. The baseline must include the resource usage associated with the server processes.

What should you do?

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

Answer: C

Explanation

sys.dm_os_memory_objects returns memory objects that are currently allocated by SQL Server. You can use sys.dm_os_memory_objects to analyze memory use and to identify possible memory leaks. Example: The following example returns the amount of memory allocated by each memory object type.

```
SELECT SUM (pages_in_bytes) as 'Bytes Used', type
FROM sys.dm_os_memory_objects
GROUP BY type
ORDER BY 'Bytes Used' DESC;
GO
```

NO.11 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:

- * The performance counters of the operating system

- * 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 SQL Server Profiler to collect performance data.

You use SQL Server Management Studio (SSMS) to analyze the data.

Does this meet the goal?

A. No

B. Yes

Answer: B

NO.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 this series. Information and details provided in a question apply only to that question.

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

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

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

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

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

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

What should you do?

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

Answer: D

Explanation

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

NO.13 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.error_reported 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_column_statistics event added.

Answer: A

Explanation

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=>

NO.14 Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a raid array. The server hosts three Microsoft SQL Server instances.

There are many SQL jobs that run during off-peak hours.

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

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

- A.** Create a Performance Monitor Data Collector Set.
- B.** Create a sys.dm_os_memory_objecs query.
- C.** Create a sys.dm_exec_sessions query.
- D.** Create a sys.dm_os_waitingtasks query.
- E.** Create a sp_configure 'max server memory' query.

Answer: C

NO.15 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

Answer: B

NO.16 Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series. You have a database that contains the following tables: BlogCategory, BlogEntry, ProductReview, Product, and SalesPerson. The tables were created using the following Transact SQL statements:

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

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

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

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

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

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

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

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

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

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

You must add the following constraints to the SalesHistory table:

- * a constraint on the SaleID column that allows the field to be used as a record identifier
- * a constant that uses the ProductID column to reference the Product column of the ProductTypes table
- * a constraint on the CategoryID column that allows one row with a null value in the column
- * a constraint that limits the SalePrice column to values greater than four Finance department users must be able to retrieve data from the 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 . 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 environment to meet the requirements for the Orders table.

What should you create?

- A.** an AFTER UPDATE trigger
- B.** a user defined function
- C.** a FOR UPDATE trigger
- D.** a stored procedure with the RETURN statement

Answer: B