

# 70-461

## Querying Microsoft SQL Server 2012

Version: Demo



## About Exambible

### *Your Partner of IT Exam*

## Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

1. You use Microsoft SQL Server 2012 to write code for a transaction that contains several statements.

There is high contention between readers and writers on several tables used by your transaction.

You need to minimize the use of the tempdb space. You also need to prevent reading queries from blocking writing queries.

Which isolation level should you use?

- A. SERIALIZABLE
- B. SNAPSHOT
- C. READ COMMITTED SNAPSHOT
- D. REPEATABLE READ

**Answer: C**

2. You develop a Microsoft SQL Server 2012 database. You create a view that performs the following tasks:

. Joins 8 tables that contain up to 500,000 records each.

. Performs aggregations on 5 fields.

The view is frequently used in several reports.

You need to improve the performance of the reports.

What should you do?

- A. Convert the view into a table-valued function.
- B. Convert the view into a Common Table Expression (CTE).
- C. Convert the view into an indexed view.
- D. Convert the view into a stored procedure and retrieve the result from the stored procedure into a temporary table.

**Answer: C**

3. You administer a Microsoft SQL Server 2012 database named ContosoDb. The database contains a table named Suppliers and a column named IsActive in the Purchases schema. You create a new user named ContosoUser in ContosoDb. ContosoUser has no permissions to the Suppliers table. You need to ensure that ContosoUser can delete rows that are not active from Suppliers. You also need to grant ContosoUser only the minimum required permissions. Which Transact-SQL statement should you use?

- A. GRANT DELETE ON Purchases.Suppliers TO ContosoUser

B. CREATE PROCEDURE Purchases.PurgeInactiveSuppliers WITH EXECUTE AS USER = 'dbo' AS DELETE FROM Purchases.Suppliers WHERE IsActive = 0 GO GRANT EXECUTE ON Purchases.PurgeInactiveSuppliers TO ContosoUser C. GRANT SELECT ON Purchases.Suppliers TO ContosoUser

D. CREATE PROCEDURE Purchases.PurgeInactiveSuppliers AS DELETE FROM Purchases.Suppliers WHERE IsActive = 0 GO GRANT EXECUTE ON Purchases.PurgeInactiveSuppliers TO ContosoUser

**Answer: D**

4. You administer a Microsoft SQL Server 2012 database that contains a table named OrderDetail. You discover that the NCI\_OrderDetail\_CustomerID non-clustered index is fragmented. You need to reduce fragmentation.

You need to achieve this goal without taking the index offline. Which Transact-SQL batch should you use?

A. CREATE INDEX NCI\_OrderDetail\_CustomerID ON OrderDetail.CustomerID WITH DROP EXISTING

B. ALTER INDEX NCI\_OrderDetail\_CustomerID ON OrderDetail.CustomerID REORGANIZE

C. ALTER INDEX ALL ON OrderDetail REBUILD

D. ALTER INDEX NCI\_OrderDetail\_CustomerID ON OrderDetail.CustomerID REBUILD

**Answer: B**

5. You are developing a database that will contain price information. You need to store the prices that include a fixed precision and a scale of six digits. Which data type should you use?

A. Float

B. Money

C. Smallmoney

D. Numeric

**Answer: D**

Explanation:

Numeric is the only one in the list that can give a fixed precision and scale.

Reference: <http://msdn.microsoft.com/en-us/library/ms179882.aspx>

6. You are developing a database application by using Microsoft SQL Server 2012.

An application that uses a database begins to run slowly.

Your investigation shows the root cause is a query against a read-only table that has a clustered index.

The query returns the following six columns:

One column in its WHERE clause contained in a non-clustered index · Four additional columns One COUNT (\*) column based on a grouping of the four additional columns

You need to optimize the statement.

What should you do?

- A. Add a HASH hint to the query.
- B. Add a LOOP hint to the query.
- C. Add a FORCESEEK hint to the query.
- D. Add an INCLUDE clause to the index.
- E. Add a FORCESCAN hint to the Attach query.
- F. Add a columnstore index to cover the query.
- G. Enable the optimize for ad hoc workloads option.
- H. Cover the unique clustered index with a columnstore index.
- I. Include a SET FORCEPLAN ON statement before you run the query.
- J. Include a SET STATISTICS PROFILE ON statement before you run the query.
- K. Include a SET STATISTICS SHOWPLAN\_XML ON statement before you run the query.
- L. Include a SET TRANSACTION ISOLATION LEVEL REPEATABLE READ statement before you run the query.
- M. Include a SET TRANSACTION ISOLATION LEVEL SNAPSHOT statement before you run the query.
- N. Include a SET TRANSACTION ISOLATION LEVEL SERIALIZABLE statement before you run the query.

**Answer: F**

7. You use Microsoft SQL Server 2012 to develop a database application. You need to create an object that meets the following requirements:

Takes an input variable

Returns a table of values

Cannot be referenced within a view

Which object should you use?

- A. Scalar-valued function
- B. Inline function
- C. User-defined data type
- D. Stored procedure

**Answer: D**

8. You administer a Microsoft SQL Server 2012 database.

The database contains a table named Employee. Part of the Employee table is shown in the exhibit. (Click the Exhibit button.)

Column name	Description
EmployeeID	<ul style="list-style-type: none"><li>• Uniquely identifies the employee record in the table</li><li>• Used throughout the database by all the other tables that reference the Employee table</li></ul>
EmployeeNum	<ul style="list-style-type: none"><li>• An alphanumeric value calculated according to company requirements</li><li>• Has to be unique within the Employee table</li><li>• Exists only within the Employee table</li></ul>
DepartmentID	<ul style="list-style-type: none"><li>• References another table named Department that contains data for each department in the company</li></ul>
ReportsToID	<ul style="list-style-type: none"><li>• Contains the EmployeeID of the manager to whom an employee reports</li></ul>

Unless stated above, no columns in the Employee table reference other tables.

Confidential information about the employees is stored in a separate table named EmployeeData. One record exists within EmployeeData for each record in the Employee table.

You need to assign the appropriate constraints and table properties to ensure data integrity and visibility.

On which column in the Employee table should you create a Primary Key constraint for this table?

- A. DateHired

- B. Departments
- C. EmployeeID
- D. EmployeeNum
- E. FirstName
- F. JobTitle
- G. LastName
- H. MiddleName
- I. ReportsToID

**Answer: C**

9. You generate a daily report according to the following query:

```
SELECT c.CustomerName
FROM Sales.Customer c
WHERE Sales.ufnGetLastOrderDate(c.CustomerID) <
    DATEADD(DAY, -90, GETDATE())
```

The Sales.ufnGetLastOrderDate user-defined function (UDF) is defined as follows:

```
CREATE FUNCTION Sales.ufnGetLastOrderDate (@CustomerID int)
RETURNS datetime
AS
BEGIN
    DECLARE @lastOrderDate datetime
    SELECT @lastOrderDate = MAX(OrderDate)
    FROM Sales.SalesOrder
    WHERE CustomerID = @CustomerID
    RETURN @lastOrderDate
END
```

You need to improve the performance of the query.

What should you do?

A. Drop the UDF and rewrite the report query as follows:

```
WITH cte(CustomerID, LastOrderDate) AS (
SELECT CustomerID, MAX(OrderDate) AS [LastOrderDate]
FROM Sales.SalesOrder
GROUP BY CustomerID
```

)

```
SELECT c.CustomerName
FROM cte
INNER JOIN Sales.Customer c
ON cte.CustomerID = c.CustomerID
WHERE cte.LastOrderDate < DATEADD(DAY, -90, GETDATE())
```

B. Drop the UDF and rewrite the report query as follows:

```
SELECT c.CustomerName
FROM Sales.Customer c
WHERE NOT EXISTS (
SELECT s.OrderDate
FROM Sales.SalesOrder
WHERE s.OrderDate > DATEADD(DAY, -90, GETDATE())
AND s.CustomerID = c.CustomerID)
```

C. Drop the UDF and rewrite the report query as follows:

```
SELECT DISTINCT c.CustomerName
FROM Sales.Customer c
INNER JOIN Sales.SalesOrder s
ON c.CustomerID = s.CustomerID
WHERE s.OrderDate < DATEADD(DAY, -90, GETDATE())
```

D. Rewrite the report query as follows:

```
SELECT c.CustomerName
FROM Sales.Customer c
WHERE NOT EXISTS (SELECT OrderDate FROM
Sales.ufnGetRecentOrders(c.CustomerID, 90))
```

Rewrite the UDF as follows:



```
CREATE FUNCTION Sales.ufnGetRecentOrders(@CustomerID int, @MaxAge datetime)
```

```
RETURNS TABLE AS RETURN (
```

```
SELECT OrderDate
```

```
FROM Sales.SalesOrder
```

```
WHERE s.CustomerID = @CustomerID
```

```
AND s.OrderDate > DATEADD(DAY, -@MaxAge, GETDATE()))
```

**Answer:** A

#### 10. DRAG DROP

You use Microsoft SQL Server 2012 to develop a database that has two tables named Div1Cust and Div2Cust.

Each table has columns named DivisionID and CustomerId . None of the rows in Div1Cust exist in Div2Cust.

You need to write a query that meets the following requirements:

The rows in Div1Cust must be combined with the rows in Div2Cust. The result set must have columns named Division and Customer. Duplicates must be retained.

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



The interface shows a list of SQL statements on the left and a drag-and-drop area on the right. The statements are:

- EXCEPT
- SELECT DivisionID, CustomerID FROM Div2Cust
- SELECT DISTINCT DivisionID, CustomerID FROM Div1Cust, Div2Cust
- INTERSECT
- SELECT DivisionID AS Division, CustomerID AS Customer FROM Div1Cust
- UNION ALL
- INNER JOIN
- UNION
- SELECT DivisionID, CustomerID FROM Div1Cust, Div2Cust
- ON Div1Cust.CustID = Div2Cust.CustID
- SELECT DivisionID, CustomerID FROM Div1Cust

The drag-and-drop area on the right is currently empty, with navigation arrows (left and right) visible.

**Answer:**

EXCEPT	
SELECT DivisionID, CustomerID FROM Div2Cust	SELECT DivisionID AS Division, CustomerID AS Customer FROM Div1Cust
SELECT DISTINCT DivisionID, CustomerID FROM Div1Cust, Div2Cust	
INTERSECT	
SELECT DivisionID AS Division, CustomerID AS Customer FROM Div1Cust	UNION ALL
UNION ALL	
INNER JOIN	
UNION	
SELECT DivisionID, CustomerID FROM Div1Cust, Div2Cust	SELECT DivisionID, CustomerID FROM Div1Cust, Div2Cust
ON Div1Cust.CustID = Div2Cust.CustID	
SELECT DivisionID, CustomerID FROM Div1Cust	

11. A table named Profits stores the total profit made each year within a territory. The Profits table has columns named Territory, Year, and Profit.

You need to create a report that displays the profits made by each territory for each year and its previous year.

Which Transact-SQL query should you use?

- A. SELECT Territory, Year, Profit, LEAD(Profit, 1, 0) OVER (PARTITION BY Territory ORDER BY Year) AS PrevProfit FROM Profits
- B. SELECT Territory, Year, Profit, LAG(Profit, 1, 0) OVER (PARTITION BY Year ORDER BY Territory) AS PrevProfit FROM Profits
- C. SELECT Territory, Year, Profit, LAG(Profit, 1, 0) OVER (PARTITION BY Territory ORDER BY Year) AS PrevProfit FROM Profits
- D. SELECT Territory, Year, Profit, LEAD(Profit, 1, 0) OVER (PARTITION BY Year ORDER BY Territory) AS PrevProfit FROM Profits

**Answer: C**

12. You use a Microsoft SQL Server 2012 database.

You want to create a table to store Microsoft Word documents.

You need to ensure that the documents must only be accessible via Transact-SQL queries.

Which Transact-SQL statement should you use?

- A. CREATE TABLE DocumentStore ( [Id] INT NOT NULL PRIMARY KEY, [Document] VARBINARY(MAX) NULL ) GO
- B. CREATE TABLE DocumentStore ( [Id] hierarchyid, [Document] NVARCHAR NOT NULL ) GO
- C. CREATE TABLE DocumentStore AS FileTable
- D. CREATE TABLE DocumentStore ( [Id] [uniqueidentifier] ROWGUIDCOL NOT NULL UNIQUE, [Document] VARBINARY(MAX) FILESTREAM NULL ) GO

**Answer: A**

13. You are a database developer at an independent software vendor. You create stored procedures that contain proprietary code.

You need to protect the code from being viewed by your customers.

Which stored procedure option should you use?

- A. ENCRYPTBYKEY
- B. ENCRYPTION
- C. ENCRYPTBYPASSPHRASE
- D. ENCRYPTBYCERT

**Answer: B**

14. You are developing a database application by using Microsoft SQL Server 2012.

An application that uses a database begins to run slowly.

You discover that during reads, the transaction experiences blocking from concurrent updates.

You need to ensure that throughout the transaction the data maintains the original version.

What should you do?

- A. Add a HASH hint to the query.
- B. Add a LOOP hint to the query.
- C. Add a FORCESEEK hint to the query.
- D. Add an INCLUDE clause to the index.
- E. Add a FORCESCAN hint to the Attach query.

- F. Add a columnstore index to cover the query.
- G. Enable the optimize for ad hoc workloads option.
- H. Cover the unique clustered index with a columnstore index.
- I. Include a SET FORCEPLAN ON statement before you run the query.
- J. Include a SET STATISTICS PROFILE ON statement before you run the query.
- K. Include a SET STATISTICS SHOWPLAN\_XML ON statement before you run the query.
- L. Include a SET TRANSACTION ISOLATION LEVEL REPEATABLE READ statement before you run the query.
- M. Include a SET TRANSACTION ISOLATION LEVEL SNAPSHOT statement before you run the query.
- N. Include a SET TRANSACTION ISOLATION LEVEL SERIALIZABLE statement before you run the query.

**Answer: M**

15. You are developing a database application by using Microsoft SQL Server 2012. You have a query that runs slower than expected.

You need to capture execution plans that will include detailed information on missing indexes recommended by the query optimizer.

What should you do?

- A. Add a HASH hint to the query.
- B. Add a LOOP hint to the query.
- C. Add a FORCESEEK hint to the query.
- D. Add an INCLUDE clause to the index.
- E. Add a FORCESCAN hint to the Attach query.
- F. Add a columnstore index to cover the query.
- G. Enable the optimize for ad hoc workloads option.
- H. Cover the unique clustered index with a columnstore index.
- I. Include a SET FORCEPLAN ON statement before you run the query.
- J. Include a SET STATISTICS PROFILE ON statement before you run the query.
- K. Include a SET STATISTICS SHOWPLAN\_XML ON statement before you run the query.



L. Include a SET TRANSACTION ISOLATION LEVEL REPEATABLE READ statement before you run the query.

M. Include a SET TRANSACTION ISOLATION LEVEL SNAPSHOT statement before you run the query.

N. Include a SET TRANSACTION ISOLATION LEVEL SERIALIZABLE statement before you run the query.

**Answer: K**

## Relate Links

**100% Pass Your 70-461 Exam with Exambible Prep Materials**

<http://www.exambible.com/70-461-exam/>

## Contact us

**We are proud of our high-quality customer service, which serves you around the clock 24/7.**

**Viste - <http://www.exambible.com/>**