Braindump2go Microsoft Exam 70-469 Updated For Free Download(1-15)

QUESTION 1You use SQL Server 2014 to maintain the data used by applications at your company.You want to execute two statements.You need to guarantee that either both statements succeed, or both statements fail together as a batch.Which code should you use?

A.   Option AB.   Option BC.   Option CD.   Option DE.   Option E Answer: DExplanation:Structure should be:BEGIN TRY BEGIN TRANSACTION..COMMIT TRANSACTIONEND TRYBEGIN CATCHROLLBACK TRANSACTIONEND CATCH. TRY...CATCH (Transact-SQL) QUESTION 2You deploy a SQL Server instance named SQLProd that uses SQL Server 2014. You need to recommend a solution to monitor the transactions that are running currently against SQLProd. The solution must minimize the amount of custom code required.What should you recommend? A.   StatisticsB.   A dynamic management viewC.   A triggerD.   User-defined views Answer: BExplanation:Dynamic management views and functions return server state information that can be used to monitor the health of a server instance, diagnose problems, and tune performance.Transactions can be monitored.Dynamic Management Views and Functions (Transact-SQL) QUESTION 3Your network contains an Active Directory domain that has two groups named Group1 and Group2. The domain contains two SQL Server instances named SQLDev and SQLProd. Each SQL Server instance has access to various storage media.The SQL Server instances have a database that contains a table named Table1. Table1 contains a column named Column1. The value for Column1 can be either Value1 or Value2. You need to recommend a solution to ensure that users in Group1 can retrieve only rows from Column1 that contain the value of Value1.What should you recommend? A.   A dynamic management viewB.   FilegroupsC.   Snapshot isolationD.   User-defined views Answer: D Explanation:A view is a virtual table whose contents are defined by a query. Like a table, a view consists of a set of named columns and rows of data. Unless indexed, a view does not exist as a stored set of data values in a database. The rows and columns of data come from tables referenced in the query defining the view and are produced dynamically when the view is referenced. A view acts as a filter on the underlying tables referenced in the view. http://msdn.microsoft.com/en-us/library/ms190174.aspx Incorrect:Not A:

Dynamic management views and functions return server state information that can be used to monitor the health of a server instance, diagnose problems, and tune performance. http://technet.microsoft.com/en-us/library/ms188754(v=sql.120).aspx QUESTION 4You have a SQL Server instance on a server named Server1. You need to recommend a solution to perform the following tasks every week: - Rebuild the indexes by using a new fill factor.- Run a custom T-SQL command.- Back up the databases. What should you recommend?More than one answer choice may achieve the goal. Select the BEST answer. A.   A triggerB.   An alertC.   A maintenance planD.   Windows PowerShellE.   A system policy Answer: CExplanation:Maintenance plans create a workflow of the tasks required to make sure that your database is optimized, regularly backed up, and free of inconsistencies.Maintenance Plans QUESTION 5Hotspot QuestionYou use SQL Server 2014. You create a table within a database by using the following DDL: CREATE TABLE Ord
(
OrderID INT IDENT
OrderDate SMALLD/
CustomerID INT,
IsTaxable INT,
SubTotal SmallMon
TaxAmount AS (Ca:
Freight SmallMone
OrderReturnedDate
OrderReturnedCust
OrderReturnedEval
)

The following table illustrates a representative sample of data:

OrderID	OrderDate	CustomerID	IsTaxable	SubTotal	TaxAmount	Freight
1	11/13/2013 11:22	58465	NULL	$ 25.99	NULL	$ 5.40
2	11/15/2013 9:34	12588	NULL	$ 42.00	NULL	NULL
3	12/1/2013 14:34	85477	NULL	$ 23.99	NULL	$ 4.85
4	12/17/2013 4:31	58742	NULL	$ 19.00	NULL	NULL
5	1/3/2014 8:22	85477	NULL	$ 12.50	NULL	$ 5.40
6	1/5/2014 18:39	63214	NULL	$ 5.69	NULL	NULL
7	1/15/2014 14:22	85471	NULL	$ 18.99	NULL	$ 7.85
8	1/19/2014 3:20	85412	NULL	$ 65.77	NULL	NULL
9	1/22/2014 13:44	12588	NULL	$ 22.38	NULL	$ 7.35
10	1/28/2014 10:14	85471	1	$ 24.99	$ 2.19	$ 5.40

The system is expected to handle 50 million orders a month over the next five years. You have been instructed by your Team Lead to follow best practices for storage and performance in the utilization of SPARSE columns.Which columns should you designate as SPARSE? To answer, mark each column as SPARSE or NOT SPARSE in the answer area. **Answer Area**

Column Names	Sparse
OrderID	O
OrderDate	O
CustomerID	
IsTaxable	O
SubTotal	O
TaxAmount	O
Freight	O

Answer:

Answer Area

Column Names	Sparse	Not Sparse
OrderID	○	◉
OrderDate	○	◉
CustomerID	○	◉
IsTaxable	◉	○
SubTotal	○	◉
TaxAmount	◉	○
Freight	◉	○

] Explanation:Note:Sparse columns are ordinary columns that have an optimized storage for null values. Sparse columns reduce the space requirements for null values at the cost of more overhead to retrieve nonnull values. Consider using sparse columns when the space saved is at least 20 percent to 40 percent.Use Sparse Columns QUESTION 6Drag and Drop QuestionYou plan to deploy a database to SQL Azure.You are designing two stored procedures named USP_1 and USP_2 that have the following requirements: Prevent data read by USP_1 from being modified by other active processes.Allow USP_2 to perform dirty reads.You need to recommend the isolation level for the stored procedures.The solution must maximize concurrency.Which isolation levels should you recommend? To answer, drag the appropriate isolation level to the correct stored procedure in the answer area.

Isolation Levels

read committed

read uncommitted

repeatable read

serializable

Answer:

Isolation Levels		Answer Area	
read committed		SP1	repeatable read
		SP2	read uncommitted
serializable			

]Explanation:Note:* SP1: repeatable reada repeatable read scan retains locks on every row it touches until the end of the transaction. Even rows that do not qualify for the query result remain locked. These locks ensure that the rows touched by the query cannot be updated or deleted by a concurrent session until the current transaction completes (whether it is committed or rolled back).* SP2: read uncommitted permits repeatable readsTransaction Isolation Levels QUESTION 7Drag and Drop QuestionYou have a SQL Azure database named Database1.You need to design the schema for a table named table1. Table1 will have less than one million rows. Table1 will contain the following information for each row:

Column	Description
ID	An incremental numeric value used to identify the row
Name	A string in English
Code	An alphanumeric code that has five characters
ModifiedDate	The date of the last modification

The solution must minimize the amount of space used to store each row. Which data types should you recommend for each column? To answer, drag the appropriate data type to the correct column in the answer area.

Data Types

int
bigint
varchar
nvarchar
char
smalldatetime
date

Answer Area

ID	Data ty
Name	Data ty
Code	Data ty
ModifiedDate	Data ty

Braindump2go.com

Answer:

Data Types

bigint
nvarchar
smalldatetime

Braindump2go.com

Answer Area

ID	int
Name	varchar
Code	char
ModifiedDate	date

] Explanation:Data Types (Transact-SQL) QUESTION 8You have a SQL Server 2014 instance named SQL1.SQL1 creates error events in the Windows Application event log.You need to recommend a solution that will run an application when SQL1 logs a specific error in the Application log.Which SQL elements should you include in the recommendation? (Each correct answer presents part of the solution. Choose all that apply.) A.   A policyB.   A maintenance planC.   An alertD.   A jobE.   A trigger Answer: DE Explanation:Use a trigger that starts a job which executes the application. http://technet.microsoft.com/en-us/library/hh849759.aspx QUESTION 9You have a database named DB1.You plan to create a stored procedure that will insert rows into three different tables. Each insert must use the same identifying value for each table, but the value must increase from one invocation of the stored procedure to the next.Occasionally, the identifying value must be reset to its initial value. You need to design a mechanism to hold the identifying values for the stored procedure to use.What should you do?More than one answer choice may achieve the goal. Select the BEST answer. A.   Create a sequence object that holds the next value in the sequence. Retrieve the next value by using the stored procedure. Reset the value by using an ALTER SEQUENCE statement as needed.B.   Create a sequence object that holds the next value in the sequence. Retrieve the next value by using the stored procedure. Increment the sequence object to the next value by using an ALTER SEQUENCE statement. Reset the value as needed by using a different ALTER SEQUENCE statement.C.   Create a fourth table that holds the next value in the sequence. At the end each transaction, update the value by using the stored procedure. Reset the value as needed by using an UPDATE statement.D.   Create an identity column in each of the three tables. Use the same seed and the same increment for each table. Insert new rows into the tables by using the stored procedure. Use the DBCC CHECKIDENT command to reset the columns as needed. Answer: AExplanation:* an application can obtain the next sequence number without inserting the row by calling the NEXT VALUE FOR function.* ALTER SEQUENCEIncludes argument:RESTART [WITH <constant>]The next value that will be returned by the sequence object. If provided, the RESTART WITH value must be an integer that is less than or equal to the maximum and greater than or equal to the minimum value of the sequence object. If the

WITH value is omitted, the sequence numbering restarts based on the original CREATE SEQUENCE options.* CREATE SEQUENCECreates a sequence object and specifies its properties. A sequence is a user-defined schema bound object that generates a sequence of numeric values according to the specification with which the sequence was created. The sequence of numeric values is generated in an ascending or descending order at a defined interval and can be configured to restart (cycle) when exhausted. Sequence Numbers QUESTION 10You have two SQL Server instances named SQLDev and SQLProd that have access to various storage media.You plan to synchronize SQLDev and SQLProd.You need to recommend a solution that meets the following requirements: - The database schemas must be synchronized from SQLDev to SQLProd. - The database on SQLDev must be deployed to SQLProd by using a package. - The package must support being deployed to Windows Azure SQL Database. What should you recommend?More than one answer choice may achieve the goal. Select the BEST answer. A.   A database snapshotB.   SQL Server Integration Services (SSIS)C.   Change data captureD.   A data-tier application Answer: BExplanation:* SIS supports connections to SQL Database by using the ADO.NET provider. OLEDB is not supported at this time. You can build the SSIS package connecting to SQL Database and create the data flow tasks the same way as you would against a typical on-premise SQL Server.SSIS for Azure and Hybrid Data Movement QUESTION 11You manage a SQL Server 2014 instance that contains a database named DB1. Users report that some queries to DB1 take longer than expected. Although most queries run in less than one second, some queries take up to 20 seconds to run. You need to view all of the performance statistics for each database file.Which method should you use? A.   Query the sys.dm_os_tasks dynamic management view.B.   Query the sys.dm_os_performance_counters dynamic management view.C.   Query the sys.dm_io_virtual_file_stats dynamic management function.D.   Examine the Data File I/O pane in Activity Monitor. Answer: CExplanation:sys.dm_io_virtual_file_statsReturns I/O statistics for data and log files. sys.dm_io_virtual_file_stats (Transact-SQL) QUESTION 12Drag and Drop QuestionYou plan to deploy SQL Server 2014.Your company identifies the following monitoring requirements for the database; An e-mail message must be sent if the SQL Server Authentication mode changes. An e-mail message must be sent if CPU utilization exceeds 90 percent. You need to identify which feature meets each monitoring requirement.Which features should you identify? To answer, drag the appropriate feature to the correct monitoring requirement in the answer area

] Explanation:* Policy-Based Management is helpful in resolving the issues presented in the following scenario:A company policy prohibits enabling Database Mail or SQL Mail. A policy is created to check the server state of those two features. An administrator compares the server state to the policy. If the server state is out of compliance, the administrator chooses the Configure mode and the policy brings the server state into compliance.* Events are generated by SQL Server and entered into the Microsoft Windows application log. SQL Server Agent reads the application log and compares events written there to alerts that you have defined. When SQL Server Agent finds a match, it fires an alert, which is an automated response to an event. In addition to monitoring SQL Server events, SQL Server Agent can also monitor performance conditions and Windows Management Instrumentation (WMI) events.To define an alert, you specify:The name of the alert.The event or performance condition that triggers the alert. The action that SQL

Server Agent takes in response to the event or performance condition. Configure Alerts to Notify Policy Administrators of Policy Failures SQL Server Agent QUESTION 13You are the new database administrator for a SQL Server 2014 instance. You conduct an assessment on the instance and determine that the auto create statistics setting on the database named DB1 has been turned off. You see no evidence that any maintenance has been occurring.You want to set up monitoring to see if query performance is being affected. You need to set up a monitoring process that will capture any cases where statistics could have been useful if they existed. What should you do? A.   Create a SQL Server Agent job to execute DBCC SHOWSTATISTICS on each of the primary key columns in the database.B.   Use the missing_column_statistics extended event.C.   Query the sys.statistics system view to see all cases where the statistics were last needed.D.   Write a query using the sys.dm_db_missing_index_group_stats DMV Joining to sys.indexes, filtering on is_hypothetical. Answer: BExplanation:The Missing Column Statistics event class indicates that column statistics that could have been useful for the optimizer are not available.By monitoring the Missing Column Statistics event class, you can determine if there are statistics missing for a column used by a query. This can cause the optimizer to choose a less efficient query plan than expected.Missing Column Statistics Event Class QUESTION 14You are designing a Windows Azure SQL Database for an order fulfillment system. You create a table named Sales.Orders with the following script.

```
CREATE TABLE Sales.Orders
(
    OrderID int IDENTITY(1,1) NOT NULL PRIMARY KEY,
    OrderDate datetimeoffset NOT NULL,
    CustomerID int NOT NULL
);
```

 Each order is tracked by using one of the following statuses: - Fulfilled- Shipped- Ordered- Received You need to design the database to ensure that that you can retrieve the following information: - The current status of an order- The previous status of an order.- The date when the status changed.- The solution must minimize storage. More than one answer choice may achieve the goal. Select the BEST answer. A.   To the Sales.Orders table, add three columns named Status, PreviousStatus and ChangeDate.Update rows as the order status changes.B.   Create a new table named Sales.OrderStatus that contains three columns named OrderID, StatusDate, and Status. Insert new rows into the table as the order status changes.C.   Implement change data capture on the Sales.Orders table.D.   To the Sales.Orders table, add three columns named FulfilledDate, ShippedDate, and ReceivedDate. Update the value of each column from null to the appropriate date as the order status changes. Answer: B QUESTION 15You create a stored procedure that retrieves all of the rows from a table named Table1. You need to recommend a solution to ensure that all of the statements in the stored procedure can be executed if another transaction is modifying rows in Table1 simultaneously.What should you recommend? A.   Snapshot isolationB.   A database snapshotC.   FilegroupsD.   Indexes Answer: AExplanation: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.Transaction Isolation Levels Braindump2go New Released Microsoft 70-469 Dumps PDF are Now For Free Download, 292 Latest Questions, Download It Right Now and Pass Your Exam 100%: http://www.braindump2go.com/70-469.html