

[Oct.-2016-NewInstant Download Braindump2go 70-465 Questions and Answers PDF 160q[NQ11-NQ17

2016/10 New 70-465: Designing Database Solutions for Microsoft SQL Server 2014 Exam Questions Updated Today!Free Instant Download 70-465 Exam Dumps (PDF & VCE) 160Q&As from Braindump2go.com Today! 100% Real Exam Questions! 100% Exam Pass Guaranteed! 1.|2016/10 Latest 70-465 Exam Dumps (PDF & VCE) 160Q&As Download:

http://www.braindump2go.com/70-465.html 2.|2016/10 Latest 70-465 Exam Questions & Answers:

https://drive.google.com/folderview?id=0B9YP8B9sF_gNVGJOa0pteHZWXM&usp=sharing QUESTION 11 You have a SQL Server 2014 environment That contains 20 servers.The corporate security policy states that all SQL Server 2014 instances must meet specific security standards.You need to recommend a management strategy for the SQL Server 2014 servers.What should you include in the recommendation? More than one answer choice may achieve the goal. Select the BEST answer. A. Multi server jobs B. Policy-Based ManagementC. Common criteria complianceD. Maintenance plans Answer: BExplanation:Policy-Based Management is a system for managing one or more instances of SQL Server. When SQL Server policy administrators use Policy-Based Management, they use SQL Server Management Studio to create policies to manage entities on the server, such as the instance of SQL Server, databases, or other SQL Server objects. QUESTION 12 Drag and Drop Question You are designing a database for a university. The database will contain two tables named Classes and Enrollment that have the following specifications:
- Classes will store brochures in the XPS format. The brochures must be structured in folders and must be accessible by using UNC paths.
- Enrollment will store information about students and their classes. Performance must be enhanced for queries of the current enrollments.You need to identify which SQL Server technology meets the specifications of each table. Which technologies should you identify? To answer, drag the appropriate technology to the correct table in the answer area.

Technologies	Answer Area
FileStream	Classes Technology
FileTable Partitioned tables	Enrollment Technology
Partitioned views	

Answer:

Technologies	Answer Area
FileStream	Classes FileTable
FileTable Partitioned tables	Enrollment Partitioned tables
Partitioned views	

Explanation:- The FileTable feature brings support for the Windows file namespace and compatibility with Windows applications to the file data stored in SQL Server. FileTable lets an application integrate its storage and data management components, and provides integrated SQL Server services-including full-text search and semantic search-over unstructured data and metadata. In other words, you can store files and documents in special tables in SQL Server called FileTables, but access them from Windows applications as if they were stored in the file system, without making any changes to your client applications. QUESTION 13 You deploy a database by using SQL Server 2014. The database contains a table named Table1.You need to recommend a solution to track all of the deletions executed on Table1. The solution must minimize the amount of custom code required.What should you recommend? A. Change data captureB. StatisticsC. A triggerD. Master Data Services Answer: AExplanation:Change data capture is designed to capture insert, update, and delete activity applied to SQL Server tables, and to make the details of the changes available in an easily consumed relational format. The change tables used by change data capture contain columns that mirror the column structure of a tracked source table, along with the metadata needed to understand the changes that have occurred. QUESTION 14 You have four databases that are accessed by using an Online Transaction Processing (OLTP) application. The databases are stored on a server named SQL1 that has SQL Server 2014 installed. You plan to deploy an additional server that has

SQL Server 2014 installed. You need to design a high-availability solution for the databases that meets the following requirements:- If SQL1 fails, the databases must be available.- Users must be able to run reports against a secondary copy of the databases.What should you include in the design? More than one answer choice may achieve the goal. Select the BEST answer. A. AlwaysOn availability groupsB. Database mirroringC. Log shippingD. Failover Clustering Answer: AExplanation:The AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012, AlwaysOn Availability Groups maximizes the availability of a set of user databases for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases, that fail over together. QUESTION 15You are troubleshooting an application that runs a query. The application frequently causes deadlocks. You need to identify which transaction causes the deadlock. What should you do? More than one answer choice may achieve the goal. Select the BEST answer. A. Query the sys.dm_exec_requests dynamic management view.B. Create a trace in SQL Server Profiler that contains the Deadlock graph event.C. Query the sys.dm_exec_sessions dynamic management view.D. Create an extended events session to capture deadlock information. Answer: DExplanation: Troubleshooting deadlocksYou have been receiving reports from users indicating that certain applications are returning deadlock errors. To maximize the effectiveness of troubleshooting these problems, you decide to focus on the deadlocks that are hit most frequently. You create an Extended Events session that:- Configures deadlock event tracking for the session.- Specifies a target that aggregates based on an identifier for the deadlock.You run the Extended Events session, and after additional deadlocks are reported you are able to obtain aggregated deadlock information, along with the complete XML deadlock graph for each source. Using this information you are able to pin point the most common deadlocks and start working on a solution. QUESTION 16Drag and Drop QuestionYou have a server named SQL1 that has SQL Server 2012 installed. SQL1 hosts a database named Database1. Database1 contains a table named Table1. Table1 is partitioned across five filegroups based on the Date field. The schema of Table1 is configured as shown in the following table.

Column	Data type
ID	Bigint
Account	Bigint
Amount	Decimal
TransactionType	Int
TransactionDate	Date

Table1 contains the indexes shown in the following table.

Index	Type	Columns
PK_Table1	Clustered, primary key	ID, TransactionDate
IX_Account	Nonclustered	Account
IX_Type	Nonclustered	TransactionType
IX_Date	Nonclustered	TransactionDate
IX_Amount	Nonclustered	Amount

You need to recommend an index strategy to maximize performance for the queries that consume the indexes available to Table1. Which type of index storage should you recommend? To answer, drag the appropriate index storage type to the correct index in the answer area.

Index Storage Types	Answer Area								
<div style="border: 1px solid gray; padding: 2px; margin-bottom: 5px;">Aligned</div> <div style="border: 1px solid gray; padding: 2px; margin-bottom: 5px;">Clustered</div>	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 30%;">IX_Type</td> <td style="width: 70%;"><div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div></td> </tr> <tr> <td>IX_Account</td> <td><div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div></td> </tr> <tr> <td>IX_DATE</td> <td><div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div></td> </tr> <tr> <td>IX_Amount</td> <td><div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div></td> </tr> </table>	IX_Type	<div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div>	IX_Account	<div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div>	IX_DATE	<div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div>	IX_Amount	<div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div>
IX_Type	<div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div>								
IX_Account	<div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div>								
IX_DATE	<div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div>								
IX_Amount	<div style="border: 1px solid gray; padding: 2px;">Index Storage Type</div>								

Answer:

Index Storage Types	Answer Area								
<div style="border: 1px solid green; padding: 2px; margin-bottom: 5px;">Aligned</div> <div style="border: 1px solid gray; padding: 2px; margin-bottom: 5px;">Clustered</div>	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 30%;">IX_Type</td> <td style="width: 70%;"><div style="border: 1px solid gray; padding: 2px;">Aligned</div></td> </tr> <tr> <td>IX_Account</td> <td><div style="border: 1px solid red; padding: 2px;">Nonaligned</div></td> </tr> <tr> <td>IX_DATE</td> <td><div style="border: 1px solid red; padding: 2px;">Nonaligned</div></td> </tr> <tr> <td>IX_Amount</td> <td><div style="border: 1px solid red; padding: 2px;">Nonaligned</div></td> </tr> </table>	IX_Type	<div style="border: 1px solid gray; padding: 2px;">Aligned</div>	IX_Account	<div style="border: 1px solid red; padding: 2px;">Nonaligned</div>	IX_DATE	<div style="border: 1px solid red; padding: 2px;">Nonaligned</div>	IX_Amount	<div style="border: 1px solid red; padding: 2px;">Nonaligned</div>
IX_Type	<div style="border: 1px solid gray; padding: 2px;">Aligned</div>								
IX_Account	<div style="border: 1px solid red; padding: 2px;">Nonaligned</div>								
IX_DATE	<div style="border: 1px solid red; padding: 2px;">Nonaligned</div>								
IX_Amount	<div style="border: 1px solid red; padding: 2px;">Nonaligned</div>								

Explanation:Index Storage Type- Designing a partitioned index independently (unaligned) of the base table can be useful in the following cases:- The base table has not been partitioned.- The index key is unique and it does not contain the partitioning column of the table. - You want the base table to participate in collocated joins with more tables using different join columns.

!!!RECOMMEND!!! 1.Braindump2go|2016/10 Latest 70-465 Exam Dumps (PDF & VCE) 160Q&As Download:
<http://www.braindump2go.com/70-465.html> 2.Braindump2go|2016/10 Latest 70-465 Exam Questions & Answers:
https://drive.google.com/folderview?id=0B9YP8B9sF_gNVGJOa0pteHZWWMXM&usp=sharing