

[May-2018-NewValid 70-764 VCE and 70-764 PDF Free Download in Braindump2go[127-137

2018 May New Microsoft 70-764 Exam Dumps with PDF and VCE Just Updated Today! Following are some new 70-764 Real Exam Questions:1. |2018 Latest 70-764 Exam Dumps (PDF & VCE) 332Q

Download:<https://www.braindump2go.com/70-764.html>2. |2018 Latest 70-764 Exam Questions & Answers

Download:<https://drive.google.com/drive/folders/0B75b5xYLjSSNdIF6dzFQVE9kUjA?usp=sharing>QUESTION 127You administer a SQL Server 2014 instance. Users report that the SQL Server has seemed slow today. A large database was being restored for much of the day, which could be causing issues. You want to write a query of the system views that will report the following:- Number of users that have a connection to the server- Whether a user's connection is active- Whether any connections are blocked- What queries are being executed- Whether the database restore is still executing and, if it is, what percentage of the restore is complete. Which system objects should you use in your query to best achieve this task?A. sys.dm_exec_requests, sys.dm_exec_sessions, sys.objectsB. sys.dm_exec_sessions, sys.dm_exec_query_stats, sys.dm_exec_query_text, sys.objectsC. sys.sysprocesses, sys.dm_exec_query_text, sys.objectsD. sys.dm_exec_requests, sys.dm_exec_sessions, sys.dm_exec_query_textAnswer: DExplanation:- sys.dm_exec_requestsReturns information about each request that is executing within SQL Server.- sys.dm_exec_sessionsReturns 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.- sys.dm_exec_query_textReturns the text of the SQL batch that is identified by the specified sql_handle.Incorrect:- sys.dm_exec_query_statsReturns aggregate performance statistics for cached query plans in SQL Server. The view contains one row per query statement within the cached plan, and the lifetime of the rows are tied to the plan itself.- sys.objectsContains a row for each user-defined, schema-scoped object that is created within a database.QUESTION 128You have a database hosted on SQL Server 2012 R2. The database contains 5 million rows. You need to recommend a repeatable method to migrate the database to SQL Azure. Which method should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.A. Create a SQL Server Integration Services (SSIS) package, and then run the package.B. Back up the database, and then restore the database.C. Extract a data-tier application, and then import the application.D. Generate scripts to create all of the all database objects and all of the data, and then execute the scripts by using SQL Azure. Answer: AExplanation:SQL Server Integration ServicesMost flexibilityData Transfer Efficiency: Good- SSIS can be used to perform a broad range of data migration tasks. SSIS provides support for complex workflow and data transformation between the source and destination. It is a good choice to transfer of data for databases that require many changes to work on Microsoft Azure SQL Database. You can use SSIS data transfer packages with another mechanism for transferring the database schema, such as a Data-tier Application package.Incorrect:Not D: Generate Scripts WizardHas explicit option for Azure SQL Database scripts generation Data Transfer Efficiency: PoorGood for smaller database- Using the Generate Scripts wizard to migrate a SQL Server database to Azure SQL Database should be limited to: Teams who have experience with the wizard.Migrating simple databases that need few schema changes to run on Azure SQL Database. The scripts generated from the source database can be modified before being used to create the new version of the database on Azure SQL Database, but using a database project in the SQL Server Data Tools has richer support for making schema changes.Migrating small databases that do not have much data. The wizard generates scripts that use insert statements instead of bulk copies to transfer the data. The insert statements can be throttled when the tables contain too much data, and are not as fast as bulk copies.QUESTION 129You have a SQL Server 2014 database named DB1. You plan to import a large number of records from a SQL Azure database to DB1. You need to recommend a solution to minimize the amount of space used in the transaction log during the import operation. What should you include in the recommendation?A. The bulk-logged recovery modelB. The full recovery modelC. A new partitioned tableD. A new log file E. A new file groupAnswer: AExplanation:Compared to the full recovery model, which fully logs all transactions, the bulk-logged recovery model minimally logs bulk operations, although fully logging other transactions. The bulk-logged recovery model protects against media failure and, for bulk operations, provides the best performance and least log space usage.Note:The bulk-logged recovery model is a special-purpose recovery model that should be used only intermittently to improve the performance of certain large-scale bulk operations, such as bulk imports of large amounts of data.QUESTION 130You have a server named Server1 that has 2 processors. You plan to deploy multiple instances of SQL Server 2014 to Server1. Each instance will have multiple databases. You need to recommend a method to allocate processor time to each database. What should you include in the recommendation? More than one answer choice may achieve the goal. Select the BEST answer.A. Resource GovernorB. Max Degree of ParallelismC. Windows System Resource Manager (WSRM)D. Processor affinityAnswer: AExplanation:SQL Server Resource

Governor is a feature that you can use to manage SQL Server workload and system resource consumption. Resource Governor enables you to specify limits on the amount of CPU, physical IO, and memory that incoming application requests can use. Incorrect: D: PROCESS AFFINITY Enables hardware threads to be associated with CPUs. QUESTION 131 You have a SQL Server 2014 environment that includes four servers. The servers are configured as shown in the following table. You plan to configure Policy-Based Management to enforce the following rules:- On Server1, enable SQL Server password policies and enable the default trace.- On Server3, ensure that the names of user-defined stored procedures begin with the prefix "usp_" and ensure that all databases use a case-sensitive collation. You need to recommend which server you must configure as a Central Management Server. Which server should you recommend? (Each correct answer presents a complete solution. Choose all that apply.) A. Server1 B. Server2 C. Server3 D. Server4 Answer: ACE Explanation: Need Standard or Enterprise edition of SQL Server. QUESTION 132 You are the administrator for a SQL Server 2014 instance that stores the data for an online transaction processing sales system. The company takes full backups every week; differential backups on the days with no full backups; and hourly transaction backups. These backups are stored on a backup server in the company's data center. Every week, the company places the full backup on a tape and sends it to a third-party backup storage system. The company is worried that a disaster might occur that could destroy their computer center and cause them to lose orders. You need to determine the best method for providing the smallest amount of data loss and downtime without leasing or purchasing additional physical locations. What should you do? More than one answer choice may achieve the goal. Select the BEST answer. A. Set up SQL Server Always On with a SQL Azure database as a replica. B. Set up SQL Server Always On by using a SQL Server on a Windows Azure Virtual Machine. C. Put the differential backup on tape and send it to the third-party backup storage system. D. Use the Microsoft SQL Server Backup to Microsoft Windows Azure Tool to direct all backups to a different geographical location. Answer: DE Explanation: SQL Server 2012 was the first version to provide the ability to back up databases to the Cloud, and SQL Server 2014 improves on the process. Microsoft SQL Server Backup to Windows Azure Tool enables backup to Windows Azure Blob Storage and encrypts and compresses SQL Server backups stored locally or in the cloud. QUESTION 133 You need to recommend a backup process for data warehouse database. The solution must meet the following requirements:- Ensure that if a hardware failure occurs, you can bring the database online without losing more than 24 hours of transactions.- Minimize the amount of administrative effort required to restore any lost data.- Minimize the space used by transaction logs. What should you include in the recommendation? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. Answer: Explanation: Box 1: Set the database to the simple recovery model. Box 2: Perform a weekly full backup of the database. Box 3: Perform a differential backup of the database every night. Note:- Simple recovery model No log backups.- Full recovery model Requires log backups QUESTION 134 You plan to deploy a database to SQL Azure. You must create two tables named Table 1 and Table 2 that will have the following specifications:- Table1 will contain a date column named Column1 that will contain a null value approximately 80 percent of the time.- Table2 will contain a column named Column2 which is the product of two other columns in Table2. Both Table1 and Table2 will contain over one million rows. You need to recommend which options must be defined for the columns. The solution must minimize the time it takes to retrieve data from the tables. Which options should you recommend? To answer, drag the appropriate options to the correct column in the answer area. Answer: Explanation:- 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.- A Persisted column would be faster to retrieve.- A computed column is computed from an expression that can use other columns in the same table. The expression can be a noncomputed column name, constant, function, and any combination of these connected by one or more operators. Unless otherwise specified, computed columns are virtual columns that are not physically stored in the table. Their values are recalculated every time they are referenced in a query. The Database Engine uses the PERSISTED keyword in the CREATE TABLE and ALTER TABLE statements to physically store computed columns in the table. Their values are updated when any columns that are part of their calculation change. QUESTION 135 You use SQL Server 2014. You create a table within a database by using the following DDL: The following table illustrates a representative sample of data: 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: Explanation: 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. QUESTION 136 You use SQL Server 2014. You need to create a single object that inserts a provided value into Table1, and then returns a count of the records in Table1. Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks. Answer: QUESTION

137 You plan to install two SQL Server 2014 environments named Environment1 and Environment2. Your company identifies the following availability requirements for each environment:- Environment1 must have mirroring with automatic failover implemented.
- Environment2 must have Always On with automatic failover implemented. You need to identify the minimum number of SQL Server 2014 servers that must be deployed to each environment to ensure that all data remains available if a physical server fails. How many servers should you identify? To answer, drag the appropriate number to the correct environment in the answer area.

Answer: **!!!!RECOMMEND!!!!**1. | 2018 Latest 70-764 Exam Dumps (PDF & VCE) 332Q

Download: <https://www.braindump2go.com/70-764.html>2. | 2018 Latest 70-764 Study Guide Video: YouTube Video:

[YouTube.com/watch?v=vq9_aKfHgXQ](https://www.youtube.com/watch?v=vq9_aKfHgXQ)