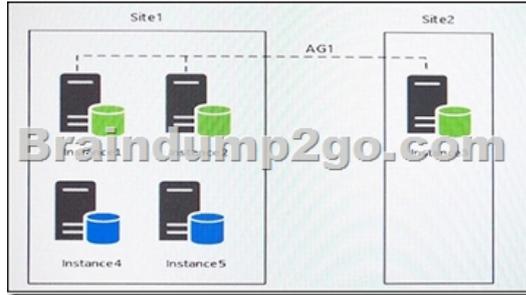


[2017-Aug-NewsHigh Quality Braindump2go 70-764 Dumps PDF 135Q Free Share[21-30

2017 August New 70-764 Exam Dumps with PDF and VCE Updated in www.Braindump2go.com Today!100% 70-764 Real Exam Questions! 100% 70-764 Exam Pass Guaranteed! 1.|2017 New 70-764 Exam Dumps (PDF & VCE) 135Q&As Download: <https://www.braindump2go.com/70-764.html> 2.|2017 New 70-764 Exam Questions and Answers Download: <https://drive.google.com/drive/folders/0B75b5xYLjSSndIF6dzFQVE9kUjA?usp=sharing> QUESTION 21Note: 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 five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram.



You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

Instance	Recovery Model
Instance1	Primary
Instance2	Synchronous
Instance3	Asynchronous

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size. Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.All databases use the full recovery model. All backups are written to the network location \SQLBackup. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

Instance	Recovery point objective
Instance 1	5 minutes
Instance 2	5 minutes
Instance 3	5 minutes
Instance 4	60 minutes
Instance 5	24 hours

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1.The wait statistics monitoring requirements for the instances are described in the following table.

Instance	Wait Statistics
Instance1	Aggregate wait statistics since the last full backup.
Instance4	Wait statistics for the current session, between session start and session end.
Instance5	Identify all the wait types for queries.

You need to create a backup plan for Instance4.Which backup plan should you create? A. Weekly full backups, nightly differential. No transaction log backups are necessary.B. Weekly full backups, nightly differential backups, transaction log backups every 5 minutes.C. Weekly full backups, nightly differential backups, transaction log backups every 12 hours.D.

Weekly full backups, nightly differential backups, nightly transaction log backups. Answer: B Explanation: From scenario: Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O. The recovery point objective of Instance4 is 60 minutes. RecoveryPoint Objectives are commonly described as the amount of data that was lost during the outage and recovery period. References: <http://sqlmag.com/blog/sql-server-recovery-time-objectives-and-recovery-point-objectives> QUESTION 22 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. A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts a customer database named DB1. Customers connect to hosted database instances by using line-of-business applications. Developers connect by using SQL Server Management Studio (SSMS). You need to grant the developers permission to alter views for DB1 while following the principle of least privilege. Which permission should you grant? A. DDLAdmin B. db_datawriter C. dbcreator D. dbo E. View Database State F. View Server State G. View Definition H. sysadmin Answer: A Explanation: To execute ALTER VIEW, at a minimum, ALTER permission on OBJECT is required. Members of the db_ddladmin fixed database role can run any Data Definition Language (DDL) command in a database. References: [https://technet.microsoft.com/en-us/library/ms190667\(v=sql.90\).aspx](https://technet.microsoft.com/en-us/library/ms190667(v=sql.90).aspx) QUESTION 23 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 observe that several indexes are fragmented. You need to rebuild the indexes. What should you use? A. Activity Monitor B. Sp_who3 C. Object Explorer in the SQL Server Management Studio (SSMS) D. SQL Server Data Collector E. SQL Server Data Tools (SSDT) F. SQL Server Configuration Manager Answer: C Explanation: How to: Rebuild an Index (SQL Server Management Studio) To rebuild an index References: [https://technet.microsoft.com/en-us/library/ms187874\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms187874(v=sql.105).aspx) QUESTION 24 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 examining information about users, sessions, and processes in an on-premises Microsoft SQL Server 2016 Standard Edition server. You need to identify waits for resources and return only the following information: What should you use? A. Activity Monitor B. Sp_who3 C. SQL Server Management Studio (SSMS) Object Explorer D. SQL Server Data Collector E. SQL Server Data Tools (SSDT) F. SQL Server Configuration Manager Answer: E Explanation: SQL Server Data Tools (SSDT) is a Microsoft Visual Studio environment for creating business intelligence solutions. SSDT features the Report Designer authoring environment, where you can open, modify, preview, save, and deploy Reporting Services paginated report definitions, shared data sources, shared datasets, and report parts. References: [https://msdn.microsoft.com/en-us/library/hh272686\(v=vs.103\).aspx](https://msdn.microsoft.com/en-us/library/hh272686(v=vs.103).aspx) QUESTION 25 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 collect performance metrics on multiple Microsoft SQL Server instances and store the data in a single repository. You need to examine disk usage, query statistics, and server activity without building custom counters. What should you use? A. Activity Monitor B. Sp_who3 stored procedure C. Object Explorer in the Microsoft SQL Server Management Studio (SSMS) D. SQL Server Data Collector E. SQL Server Data Tools (SSDT) F. SQL Server Configuration Manager Answer: D Explanation: The data collector is a core component of the data collection platform for SQL Server 2017 and the tools that are provided by SQL Server. The data collector provides one central point for data collection across your database servers and applications. This collection point can obtain data from a variety of sources and is not limited to performance data QUESTION 26 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 question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases. One customer reports that their database is not responding as quickly as the service level agreements dictate. You observe that the database is fragmented. You need to optimize query performance. Solution: You reorganize all indexes. Does the solution meet the goal? A. Yes B. No Answer: A Explanation: You can remedy index fragmentation by either reorganizing an index or by rebuilding an index. References: [https://msdn.microsoft.com/en-us/library/ms189858\(v=sql.105\).aspx](https://msdn.microsoft.com/en-us/library/ms189858(v=sql.105).aspx) QUESTION 27 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 question sets might have more than one correct solution, while others might not have a correct solution. After you answer a

question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. A company has a server that runs Microsoft SQL Server 2016 Web edition. The server has a default instance that hosts a database named DB1. You need to ensure that you can perform auditing at the database level for DB1. Solution: You migrate DB1 to the default instance on a server that runs Microsoft SQL Server 2016 Standard edition. Does the solution meet the goal? A. Yes B. No Answer: B Explanation: All editions of SQL Server support server level audits. All editions support database level audits beginning with SQL Server 2016 SP1. Prior to that, database level auditing was limited to Enterprise, Developer, and Evaluation editions. References: <https://docs.microsoft.com/en-us/sql/relational-databases/security/auditing/sql-server-audit-database-engine> QUESTION 28 You have a database that stores information for a shipping company. You create a table named Customers by running the following Transact-SQL statement. (Line numbers are included for reference only.)

```
01 CREATE TABLE dbo.Customers (  
02     customerId int,  
03     customerName varchar(200),  
04     salesPerson varchar(20)  
05 )  
06 AS  
07 RETURN SELECT 1 AS [fn_securityPredicateOrder_result]  
08 FROM dbo.Customers  
09 WHERE @salesPerson = user_name()
```

You need to ensure that salespeople can view data only for the customers that are assigned to them. Which Transact-SQL segment should you insert at line 07? A. RETURNS varchar(20) WITH SCHEMABINDING B. RETURNS dbo.Customers ORDER BY @salesPerson C. RETURNS table ORDER BY @salesPerson D. RETURNS table WITH SCHEMABINDING Answer: D Explanation: The return value can either be a scalar (single) value or a table. SELECT 1 just selects a 1 for every row, of course. What it's used for in this case is testing whether any rows exist that match the criteria: if a row exists that matches the WHERE clause, then it returns 1, otherwise it returns nothing. Specify the WITH SCHEMABINDING clause when you are creating the function. This ensures that the objects referenced in the function definition cannot be modified unless the function is also modified. References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-function-transact-sql> QUESTION 29

Hotspot Question You manage a Microsoft-SQL Server database named sales Orders. You need to verify the integrity of the database and attempt to repair any errors that are found. Repair must not cause any data to be lost in the database. How should you complete the DBCC command? To answer, select the appropriate options in the answer area. Answer: Explanation: Box 1: CHECKDB DBCC CHECKDB checks the logical and physical integrity of all the objects in the specified database. Partial syntax: DBCC CHECKDB (database_name | database_id | [, NOINDEX] , { REPAIR_ALLOW_DATA_LOSS | REPAIR_FAST | REPAIR_REBUILD }) ... Box 2: REPAIR_REBUILD DBCC CHECKDB ... REPAIR_ALLOW_DATA_LOSS | REPAIR_FAST | REPAIR_REBUILD specifies that DBCC CHECKDB repair the found errors. REPAIR_REBUILD performs repairs that have no possibility of data loss. This can include quick repairs, such as repairing missing rows in non-clustered indexes, and more time-consuming repairs, such as rebuilding an index. References: <https://docs.microsoft.com/en-us/sql/t-sql/database-consolecommands/dbcc-checkdb-transact-sql> QUESTION 30

Hotspot Question You have a Microsoft SQL Server instance that hosts a database named DB1 that contains 800 gigabyte (GB) of data. The database is used 24 hours each day. You implement indexes and set the value of the Auto Update Statistics option set to True. Users report that queries take a long time to complete. You need to identify statistics that have not been updated for a week for tables where more than 1,000 rows changed. How should you complete the Transact-SQL statement? To answer, configure the appropriate Transact-SQL segments in the answer area. Answer: Explanation: Box 1: stats_date See example below. Box 2: rowmodctr See example below. Box 3: stats_date You need to identify statistics that have not been updated for a week. Box 4: rowmodctr You need to identify that more than 1,000 rows changed. Rowmodctr counts the total number of inserted, deleted, or updated rows since the last time statistics were updated for the table. Example: We will query every statistics object which was not updated in the last day and has rows modified since the last update. We will use the rowmodctr field of sys.sysindexes because it shows how many rows were inserted, updated or deleted since the last update occurred. Please note that it is not always 100% accurate in SQL Server 2005 and later, but it can be used to check if any rows were modified.--Get the list of outdated statistics

```
SELECT OBJECT_NAME(id), name, STATS_DATE(id, indid), rowmodctr FROM sys.sysindexes WHERE STATS_DATE (id, indid) <= DATEADD(DAY, -1, GETDATE()) AND rowmodctr > 0 AND id IN (SELECT object_id FROM sys.tables) GO
```

After collecting this information, we can decide which statistics require an update. !!!RECOMMEND!!! 1. | 2017 New 70-764 Exam

Dumps (PDF & VCE) 135Q&As Download: <https://www.braindump2go.com/70-764.html> 2. | 2017 New 70-764 Study Guide

Download: YouTube Video: [YouTube.com/watch?v=9L8X7qjAeJU](https://www.youtube.com/watch?v=9L8X7qjAeJU)