

[2017-Aug-News70-764 Exam VCE Dumps Free Download in Braindump2go[41-45]

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 41Note: 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.Hotspot QuestionYou are a database administrator for a company that has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases, and each customer uses a dedicated instance. The environments that you manage are shown in the following table.

Customer	Cloud Type	Description
AdventureWorks Cycles	Private	The environment includes a database named Adventureworks that contains a single schema named ADVSchema . You must implement auditing for all objects in the ADVSchema schema. You must also implement auditing to record access to data that is considered sensitive by the company.
Tailspin Toys	Private	Tailspin Toys has a custom application that accesses a hosted database named TSpinDB . The application will monitor TSpinDB and capture information over time about which database objects are accessed and how frequently they are accessed.
Contoso, Ltd.	Private	The environment has a database named ConDB that was recently upgraded to Microsoft SQL Server 2016. Contoso reports that ConDB is slow to return results when the server is busy. You must implement a process to improve the performance of the database.
Wingtip Toys	Private	Wingtip Toys has a database named WingDB . All tables in the database have indexes. Users report system response time is slow during peak activity periods. You observe that the performance issues are related to locking. Wingtip Toys receives data updates from suppliers each week. You must implement a process for importing the data into WingDB . You must use minimal logging and minimized data loss during import process.
Wide World Importers	Public	The environment includes a database named WDWDB . Neither auditing nor statistics are configured for WDWDB . You must log any deletion of views and all database record update operations.

You need to configure auditing for the Adventure Works environment.How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.NOTE: Each correct selection is worth one point.

Answer Area

USE master
GO

▼ AuditADUAccess

CREATE DATABASE AUDIT
ALTER DATABASE AUDIT
CREATE SERVER AUDIT
ALTER SERVER AUDIT

TO FILE (FILEPATH = 'C:\ADUAudit\')
WHERE object_name = 'SensitiveData'

GO

▼ AuditADUAccess WITH

CREATE DATABASE AUDIT
ALTER DATABASE AUDIT
CREATE SERVER AUDIT
ALTER SERVER AUDIT

GO

Use Adventureworks

▼ SPECIFICATION [Filter]

CREATE DATABASE AUDIT
ALTER DATABASE AUDIT
CREATE SERVER AUDIT
ALTER SERVER AUDIT

▼ [AuditADUAccess]

FOR SERVER AUDIT
FOR DATABASE AUDIT
USE [AuditDataAcces]
SELECT ID

ADD (SELECT ON SCHEMA::[ADVSchema] BY [pu
WITH (STATE = ON)
GO

Answer:

Output as PDF file has been powered by [[Universal Post Manager](#)] plugin from [www.ProfProjects.com](#) | Page 1/6 |

Answer Area

USE master
GO
▼ AuditADUAccess
CREATE DATABASE AUDIT
ALTER DATABASE AUDIT
CREATE SERVER AUDIT
ALTER SERVER AUDIT

TO FILE (FILEPATH = 'C:\ADUAudit\')
WHERE object_name = 'SensitiveData'
GO
▼ AuditADUAccess WITH (STATE = ON)
CREATE DATABASE AUDIT
ALTER DATABASE AUDIT
CREATE SERVER AUDIT
ALTER SERVER AUDIT

GO
Use Adventureworks
▼ SPECIFICATION [FilterForSensitiveData]
CREATE DATABASE AUDIT
ALTER DATABASE AUDIT
CREATE SERVER AUDIT
ALTER SERVER AUDIT

▼ [AuditADUAccess]
FOR SERVER AUDIT
FOR DATABASE AUDIT
USE [AuditDataAccess]
SELECT ID

ADD (SELECT ON SCHEMA::[ADUSchema] BY [public])
WITH (STATE = ON)
GO

Explanation:Box 1: CREATE SERVER AUDITCreate the server audit.You must implement auditing to record access to data that is considered sensitive by the company.Create database auditBox 2: ALTER SERVER AUDITEnable the server audit.Box 3: CREATE DATABASE AUDITCreate the database audit specification.Box 4: FOR SERVER AUDITIYou must implement auditing for all objects in the ADVSchema. QUESTION 42Hotspot QuestionA company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instanced. The environments host several customer databases.You host a local database and a Stretch database that has a table named Members for one specific customer.You need to provide the customer with information about the space used in the databases.In the table below, identify the query that provides the required information for each database.NOTE: Make only one selection in each column.

Answer Area

Query	Local database	Stretch database
EXEC sp_spaceused N'Company.Members'	<input type="radio"/>	<input type="radio"/>
EXEC sp_spaceused N'Company.Members', @mode = 'REMOTE_ONLY'	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Query	Local database	Stretch database
EXEC sp_spaceused N'Company.Members'	<input checked="" type="radio"/>	<input type="radio"/>
EXEC sp_spaceused N'Company.Members', @mode = 'REMOTE_ONLY'	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:Local database: EXEC sp_spaceused N'Company.Members' QUESTION 43Hotspot QuestionYou manage a Microsoft SQL Server instance. You have a user named User1.You need to grant the minimum permissions necessary to allow User1 to review audit logs.For each action, which option should you use? To answer, select the appropriate options in the answer area.

Answer Area

Actions
User1 server role assignm

Braindump2go.com

Transact-SQL syntax

Answer:

Answer Area

Actions

Options

User1 server role assignment

Transact-SQL syntax

diskadmin

serveradmin

securityadmin

sysadmin

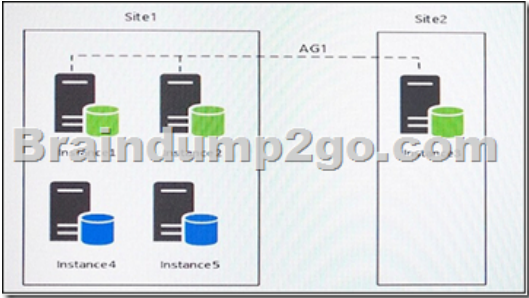
sys.server_file_audits

sys.server_audit_specifications

sys.server_file_permissions

sys.server_principals

Explanation:Box 1: securityadminTo access log files for instances of SQL Server that are online, this requires membership in the securityadmin fixed server role.Box 2: sys.server_audit_specificationssys.server_audit_specifications contains information about the server audit specifications in a SQL Server audit on a server instance. QUESTION 44Note: 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.Drag and Drop QuestionYou 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	Role
Instance1	Primary
Instance2	Synchronous secondary
Instance3	Asynchronous secondary

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 backup
Instance4	Wait statistics for the last backup
Instance5	Identify all the wait types for queries

You need to propose a new process for the StagedExternal database. Which five actions should you recommended be performed in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions


- Drop all nonclustered indexes on the target table.
- Create a transaction log backup.
- Change the recovery model of StagedExternal to SIMPLE.
- Run the nightly import process.
- Change the recovery model of StagedExternal to BULK_LOGGED.
- Drop all clustered and nonclustered indexes on the target table.
- Recreate any dropped indexes on the target table.
- Create a transaction log backup.
- Change the recovery model of StagedExternal to BULK_LOGGED.

Answer:

Actions

Answer Area

Explanation: From scenario: 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. All databases use the full recovery model. QUESTION 45 Note: 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. Drag and Drop Question 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	Role
Instance1	Primary
Instance2	Synchronous Secondary
Instance3	Asynchronous Secondary

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	Requirement
Instance1	Aggregate wait statistics since the last full backup of the database, between session terminate.
Instance4	Identify all the wait types for queries.

You need to analyze the wait type and statistics for specific instanced in the environment. Which object should you use to gather information about each instance? To answer, drag the appropriate objects to the correct instances. Each object may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.

Objects

Sys.dm_os_wait_stats

Sys.dm_exec_connections

Sys.dm_exec_requests

Sys.dm_exec_sessions

Sys.dm_exec_query_stats

Sys.dm_exec_query_resource_semaphores

Sys.dm_exec_session_wait_stats

Answer Area

Instance	Object
Instance1	<div>Object</div>
Instance4	<div>Object</div>

Answer:

Objects

Sys.dm_exec_connections

Sys.dm_exec_requests

Sys.dm_exec_sessions

Sys.dm_exec_query_resource_semaphores

Sys.dm_exec_session_wait_stats

Answer Area

Instance	Object
Instance1	<div>Sys.dm_exec_query_stats</div>
Instance4	<div>Sys.dm_os_wait_stats</div>

Explanation: Instance 1: sys.dm_exec_query_stats From Scenario: Instance1 requirement: Aggregate statistics since last server restart. sys.dm_exec_query_stats returns aggregate performance statistics for cached query plans in SQL Server. Instance 4: sys.dm_os_wait_stats sys.dm_os_wait_stats returns information about all the waits encountered by threads that executed. From

Scenario: Instance4 requirement: Identify the most prominent wait types. Instance 5:sys.dm_exec_session_wait_statsFrom Scenario:
Instance5 requirement: Identify all wait types for queries currently running on the server.sys.dm_exec_session_wait_stats returns
information about all the waits encountered by threads that executed for each session. !!!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)