

[2017-Aug-News100% Real 70-764 PDF Dumps 135Q-Braindump2go[31-40

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 31Hotspot QuestionYou are configuring log shipping for a Microsoft SQL Server database named salesOrders.You run the following Transact-SQL script:

```
DECLARE @IS_BackupJob
DECLARE @IS_Primary
DECLARE @SF_Add_RetC
EXEC @SF_Add_RetCode
@database = N's
,@backup_direct
,@backup_share
,@backup_job_na
,@backup_retent
,@backup_compre
,@backup_thresh
,@threshold_sle
,@history_reten
,@backup_job_id
,@primary_id =
,@overwrite = 1
IF (@@ERROR = 0 AND
BEGIN
DECLARE @I
DECLARE @I
,@ena
,@fre
,@fre
,@fre
,@fre
,@fre
,@fre
,@act
,@act
,@act
,@act
,@sch
,@sch
,@sch
EXEC msdb.
@job
,@sch
EXEC msdb.
@job_
,@ena
END
EXEC master.dbo.sp_a
```

You need to determine the changes that the script has on the environment.How does the script affect the environment? To answer, select the appropriate options in the answer area.NOTE: Each correct selection is worth one point.You are configuring log shipping for a Microsoft SQL Server database.

Answer Area

A dedicated file share [answer choice] used to store the backups.

A SQL Server monitor instance [answer choice] on a server named ADATUM-SQL11.

Backup must wait be deleted after [answer choice].

The backup job will run every [answer choice].

Answer: **Answer Area**

A dedicated file share [answer choice] used to store the backups.

A SQL Server monitor instance [answer choice] on a server named ADATUM-SQL11.

Backup must wait be deleted after [answer choice].


The backup job will run every [answer choice].

Explanation:Box 1: isThe dedicated backup file share is \localhostBackupBox 2: does not runThe only thing with a name related to ADATM-SQL11 is the schedule name.Box 3: 72 hours4320 minutes equals 72 hours.Note: @backup_retention_period=] backup_retention_period Isthe length of time, in minutes, to retain the log backup file in the backup directory on the primary server. backup_retention_period is int, with no default, and cannot be NULL.Box 4: 15 minutes.[@freq_subday_type =] freq_subday_type Specifies the units for freq_subday_interval. freq_subday_typeis int, with a default of 0, and can be one of these values.Here it is 4,

which means minutes. [@freq_subday_interval =] freq_subday_interval The number of freq_subday_type periods to occur between each execution of a job. freq_subday_interval is int, with a default of 0. Note: Interval should be longer than 10 seconds. freq_subday_interval is ignored in those cases where freq_subday_type is equal to 1. Here it is 15. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-schedule-transact-sql>
<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-log-shipping-primary-database-transact-sql>

QUESTION 32 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. **Hotspot 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 session, between session start and session end.
Instance5	Identify all the wait types for queries.

You need to create the connection strings for the operations and reporting systems. In the table below, identify the option that must be specified in each connection string. NOTE: Make only one selection in each column.

Answer Area	Option	Reporting system	Operations system
Connect to a Listener using ApplicationIntent=ReadOnly.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Connect to the current primary replica SQL instance.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Connect to any current read-only replica SQL instance.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Connect to a Listener.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Connect to the current primary replica SQL instance.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Answer:

Option	Reporting system	Operations system
Connect to a Listener using ApplicationIntent=ReadOnly.	<input type="radio"/>	<input type="radio"/>
Connect to the current primary replica SQL instance.	<input type="radio"/>	<input checked="" type="radio"/>
Connect to any current read-only replica SQL instance.	<input checked="" type="radio"/>	<input type="radio"/>
Connect to a Listener.	<input type="radio"/>	<input type="radio"/>
Connect to the current primary replica SQL instance.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation: Reporting system: Connect to any current read-only replica instance We configure Read-OnlyAccess on an Availability Replica. We select Read-intent only. Only read-only connections are allowed to secondary databases of this replica. The secondary database(s) are all available for read access. From Scenario: Reporting system: This solution accesses data inDB1 with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas. Operating system: Connect to the current primary replica SQL instance By default both read-write and read-intent access are allowed to the primary replica and no connections are allowed to secondary replicas of an Always On availability group. From scenario: Operations system: This solution accesses data inDB1 with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations. References:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/configure-read-only-access-on-an-availability-replica-sql-server>

QUESTION 33 Drag and Drop Question You are the database administrator for a Microsoft SQL Server instance. You develop an Extended Events package to look for events related to application performance. You need to change the event session to include SQL Server errors that are greater than error severity 15. Which five Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.


Transact-SQL segments	Answer Area
WHERE ((sqlserver.data-base_id>(4)) AND (severity>(15)))	
(ACTION(sqlserver.client_app_name, sqlserver.data-base_id, sqlserver.session_id)	
ALTER EVENT SESSION Contoso1 ON SERVER	
GO	
ADD EVENT sqlserver.error_reported	
ADD TARGET sqlserver.error_reported	

Answer:

Transact-SQL segments	Answer Area
	ALTER EVENT SESSION Contoso1 ON SERVER
	ADD EVENT sqlserver.error_reported
	(ACTION(sqlserver.client_app_name, sqlserver.data-base_id, sqlserver.session_id)
	WHERE ((sqlserver.data-base_id>(4)) AND (severity>(15)))
)
	GO
ADD TARGET sqlserver.error_reported	

Explanation: Step 1: ALTER EVENT SESSION Contoso1 ON SERVER Step 2: ADD EVENT ... Step 3: (ACTION ... Step 4: WHERE ... Step 5:) GO Example: To start an Extended Events sessions in order to trap SQL Server errors with severity greater than 10, just run the following script: CREATE EVENT SESSION [error_trap] ON SERVER ADD EVENT sqlserver.error_reported (ACTION(package0.collect_system_time, package0.last_error, sqlserver.client_app_name, sqlserver.client_hostname, sqlserver.database_id, sqlserver.database_name, sqlserver.nt_username, sqlserver.plan_handle, sqlserver.query_hash, sqlserver.session_id, sqlserver.sql_text, sqlserver.tsqf_frame, sqlserver.tsqf_stack, sqlserver.username) WHERE ((severity)>10)) ADD TARGET package0.event_file (SET

filename=N'D:Program Files\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQLXE\eventerror_trap.xml')WITH(STARTUP_STATE=OFF)GO QUESTION 34 Hotspot Question You manage a Microsoft SQL Server environment. A server fails and writes the following event to the application event log: MSG_AUDIT_FORCED_SHUTDOWN You configure the SQL Server startup parameters as shown in the following graphic:



Use the drop-down menus to select the answer choice that answers each question. NOTE: Each correct selection is worth one point.

Answer: **Answer Area**

In which user mode will the SQL Server instance start?

With which server role can a local Windows administrator connect to the database?

Explanation: Box 1: single-user The startup option -m starts an instance of SQL Server in single-user mode. Box 2: sysadmin Starting SQL Server in single-user mode enables any member of the computer's local Administrators group to connect to the instance of SQL Server as a member of the sysadmin fixed server role. References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/database-engine-service-startup-options> QUESTION 35

Drag and Drop Question You administer a database that is used for reporting purposes. The database has a large fact table that contains three hundred million rows. The table includes a clustered columnstore index and a nonclustered index on the ProductID column. New rows are inserted into the table every day. Performance of queries that filter the Product ID column have degraded significantly. You need to improve the performance of the queries. Which three actions should you perform 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	Answer Area
Drop the clustered columnstore index.	
Create a nonclustered index on ProductID.	
Drop and recreate the clustered columnstore index.	
Create a clustered rowstore index on ProductID.	
Recreate the clustered columnstore index using DROP EXISTING.	
Create a clustered rowstore index on ProductID.	
Rebuild the clustered columnstore index.	

Answer:

Actions	Answer Area
	Drop the clustered columnstore index.
	Create a clustered rowstore index on ProductID.
Drop and recreate the clustered columnstore index.	Create a nonclustered index on ProductID.
Recreate the clustered columnstore index using DROP EXISTING.	
Rebuild the clustered columnstore index.	

Explanation: Step 1: Drop the clustered columnstore index
Step 2: Create a clustered rowstore index on ProductID. Rowstore indexes perform best on queries that seek into the data, searching for a particular value, or for queries on a small range of values. Use rowstore indexes with transactional workloads since they tend to require mostly table seeks instead of table scans.
Step 3: Create a nonclustered index on ProductID !!!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)