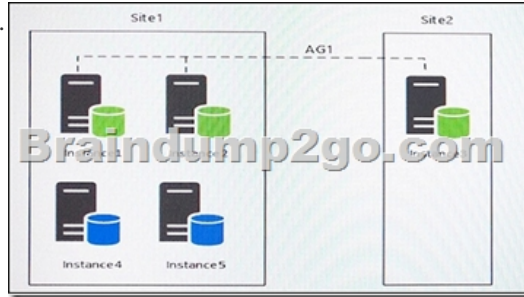


## [2017-Aug-News70-764 Dumps VCE and PDF(Full Version)135Q Download in Braindump2go[46-50

2017 August New 70-764 Exam Dumps with PDF and VCE Updated in [www.Braindump2go.com](http://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 46Note: 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  | Secondary    |
|-----------|--------------|
| 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 of the database.           |
| Instance4 | Wait statistics for the current session, between session start and session end. |
| Instance5 | Identify all the wait types for the current session.                            |

You need to configure a new replica of AG1 on Instance6.How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL statements to the correct locations. Each Transact-SQL segment 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.

**Transact-SQL segments**

DATABASE  
REPLICA  
SYNCHRONOUS\_COMMIT  
ASYNCHRONOUS\_COMMIT  
PRIMARY  
MANUAL  
SECONDARY\_ONLY

\*\*\*\*

**Answer Area**

```
ALTER AVAILABILITY GROUP AG_1 MODIFY Transact-SQL segment ON 'INSTANCE6'  
WITH (AVAILABILITY_MODE = Transact-SQL segment );  
ALTER AVAILABILITY GROUP AG_1 MODIFY Transact-SQL segment ON 'INSTANCE6'  
WITH (FAILOVER_MODE = Transact-SQL segment );
```

Answer: **Transact-SQL segments**

DATABASE  
REPLICA  
SYNCHRONOUS\_COMMIT  
ASYNCHRONOUS\_COMMIT  
PRIMARY  
MANUAL  
SECONDARY\_ONLY

\*\*\*\*

**Answer Area**

```
ALTER AVAILABILITY GROUP AG_1 MODIFY REPLICA ON 'INSTANCE6'  
WITH (AVAILABILITY_MODE = SYNCHRONOUS_COMMIT );  
ALTER AVAILABILITY GROUP AG_1 MODIFY REPLICA ON 'INSTANCE6'  
WITH (FAILOVER_MODE = AUTOMATIC );
```

Explanation:Scenario: 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.Box 1: REPLICAMODIFY REPLICA ON modifies any of the replicas of the availability group.Box 2: SYNCHRONOUS\_COMMIT You must minimize latency between the nodes in AG1 AVAILABILITY\_MODE = { SYNCHRONOUS\_COMMIT | ASYNCHRONOUS\_COMMIT } Specifies whether the primary replica has to wait for the secondary availability group to acknowledge the hardening (writing) of the log records to disk before the primary replica can commit the transaction on a given primary database.FAILOVER AUTOMATIC (box 4) requires SYNCHRONOUS\_COMMITBox 3: REPLICAMODIFY REPLICA ON modifies any of the replicas of the availability group.Box 4: AUTOMATIC You must minimize latency between the nodes in AG1 FAILOVER\_MODE = { AUTOMATIC | MANUAL } Specifies the failover mode of the availability replica that you are defining. FAILOVER\_MODE is required in the ADD REPLICA ON clause and optional in the MODIFY REPLICA ON clause.AUTOMATIC enables automatic failover. AUTOMATIC is supported only if you also specify AVAILABILITY\_MODE = SYNCHRONOUS\_COMMIT. QUESTION 47Drag and Drop QuestionYou have a test server that contains a database named DB1. Backups of the database are written to a single backup device. The backup device has a full, differential, and transaction log backup.You discover that the database is damaged. You restore the database to the point at which the differential backup was taken.You need to rebuild the database with data stored in the latest transaction logs. How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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.

**Transact-SQL statements**

NORECOVERY  
 RECOVERY  
 LOG  
 DBCC CHECKDB  
 CONTINUE\_AFTER\_ERROR  
 RESTORE V\_LIFTONLINE

.....

**Answer Area**

Transact-SQL segment DB1 FROM DISK = N'Z:\Backups\Backup.bak WITH

Transact-SQL segment

Answer: **Transact-SQL statements**

NORECOVERY  
 LOG  
 DBCC CHECKDB  
 CONTINUE\_AFTER\_ERROR  
 RESTORE V\_LIFTONLINE

.....

**Answer Area**

RESTORE DB1 FROM DISK = N'Z:\Backups\Backup.bak WITH

RECOVERY

Explanation: Box 1: RESTORE Box 2: RECOVERY The RESTORE ... WITH RECOVERY option puts the database into a useable state, so users can access a restored database. References:

<https://www.mssqltips.com/sqlservertutorial/112/recovering-a-database-that-is-in-the-restoring-state/> QUESTION 48 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 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 <b>Adventureworks</b> that contains a single schema named <b>ADVSchem</b> . You must implement auditing for all objects in the <b>ADVSchem</b> schema. You must also implement auditing to record access to data that is considered sensitive by the company.   |
| Tallspin Toys         | Private    | Tallspin Toys has a custom application that accesses a hosted database named <b>TSpinDB</b> . The application will monitor <b>TSpinDB</b> 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 <b>ConDB</b> that was recently upgraded to Microsoft SQL Server 2016. Contoso reports that <b>ConDB</b> is slow to return results when the server is busy. You must optimize performance.  |
| Wingtip Toys          | Private    | Wingtip Toys has a database named <b>WingDB</b> . 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.<br><br>Wingtip Toys receives data updates from suppliers each week. You must implement a process for importing the data into <b>WingDB</b> . You must use minimal logging and minimized data loss during import process. |
| Wide World Importers  | Public     | The environment includes a database named <b>WDWDB</b> . Neither auditing nor statistics are configured for <b>WDWDB</b> . You must log any deletion of views and all database record update operations.  |

You need to configure auditing for **WDWDB**. In the table below, identify the event type that you must audit for each activity. **Answer Area**

| Event type                              |
|---|
| <input type="checkbox"/> Data changes   |
| <input type="checkbox"/> Schema changes |
| <input type="checkbox"/> SQL batch      |
| <input type="checkbox"/> Data access    |

Answer: **Answer Area**

| Event type     | View deletions                   | Update operations                |
|----------------|----------------------------------|----------------------------------|
| Data changes   | <input type="radio"/>            | <input checked="" type="radio"/> |
| Schema changes | <input checked="" type="radio"/> | <input type="radio"/>            |
| SQL batch      | <input type="radio"/>            | <input type="radio"/>            |
| Data access    | <input type="radio"/>            | <input type="radio"/>            |

**QUESTION 49** 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 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 |  |
|-----------------------|------------|--|
| AdventureWorks Cycles | Private    | The environment named Adventureworks for all objects in the database must also implement data that is consistent.  |
| Tailspin Toys         | Private    | Tailspin Toys has a hosted application that will need information over objects are accessed.   |
| Contoso, Ltd.         | Private    | The environment that was recently Server 2016. Contoso to return results for the data. Sp... rize per... .   |
| Wingtip Toys          | Private    | Wingtip Toys has tables in the data system response periods. You observe are related to local...<br>Wingtip Toys rec... each week. You r... importing the da... minimal logging... import process. |
| Wide World Importers  | Public     | The environment named WDWDB. Neither configured for W... of views and all d... operations.   |

You need to configure the Contoso instance. How should you complete the Transact-SQL statement? To answer, select the appropriate Transact-SQL segments in the answer area.

**Answer Area**

```

sp_configure [dropdown], 1
    'max worker threads'
    'show advanced options'
    'cost threshold for parallelism'
    'max degree of parallelism'

GO
RECONFIGURE WITH OVERRIDE
sp_configure [dropdown], 0
    'max worker threads'
    'show advanced options'
    'cost threshold for parallelism'
    'max degree of parallelism'

GO
RECONFIGURE
GO
    
```

Answer: **Answer Area**

```

sp_configure [dropdown], 1
    'max worker threads'
    'show advanced options'
    'cost threshold for parallelism'
    'max degree of parallelism'

GO
RECONFIGURE WITH OVERRIDE
sp_configure [dropdown], 0
    'max worker threads'
    'show advanced options'
    'cost threshold for parallelism'
    'max degree of parallelism'

GO
RECONFIGURE
GO
    
```

Explanation:Box 1: show advanced optionsAdvanced configuration options are displayed by first setting show advanced option to 1.Box 2: max worker threadsSQL Server uses the native thread services of the operating systems so that one or more threads support each network that SQL Server supports simultaneously, another thread handles database checkpoints, and a pool of threads handles all users. The default value for max worker threads is 0. This enables SQL Server to automatically configure the number of worker threads at startup. The default setting is best for most systems.References:

[https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-the-max-worker-threads-server-configuration-op](https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-the-max-worker-threads-server-configuration-option)

tion QUESTION 50Hotspot QuestionYou are the database administrator of a Microsoft SQL Server instance. Developers are writing stored procedures to send emails using sp\_send\_dbmail. Database Mail is enabled.You need to configure each account's profile security and meet the following requirements:In the table below, identify the profile type that must be used for each account.NOTE: Make only one selection in each column.

| Profile type    | SMTP1_Account         | SMTP2_Account         |
|-----------------|-----------------------|-----------------------|
| Public Profile  | <input type="radio"/> | <input type="radio"/> |
| Default Profile | <input type="radio"/> | <input type="radio"/> |

Answer:

| Profile type    | SMTP1_Account         | SMTP2_Account                    |
|-----------------|-----------------------|----------------------------------|
| Public Profile  | <input type="radio"/> | <input type="radio"/>            |
| Default Profile | <input type="radio"/> | <input checked="" type="radio"/> |

Explanation:SMTP1\_Account1: Private ProfileWhen no profile\_name is specified, sp\_send\_dbmail uses the default private profile for the current user. If the user does not have a default private profile, sp\_send\_dbmail uses the default public profile for the msdb database.SMTP1\_Account2: Default ProfileExecute permissions forsp\_send\_dbmail default to all members of the DatabaseMailUser database role in the msdb database.References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-send-dbmail-transact-sql>

!!!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)