

[March-2019-NewExam 70-765 PDF and VCE Dumps 192Q Free Offered by Braindump2go

2019/March Braindump2go 70-765 Exam Dumps with PDF and VCE New Updated Today! Following are some new 70-765 Real Exam Questions:]1.[2019 Latest 70-765 Exam Dumps (PDF & VCE) Instant Download:

<https://www.braindump2go.com/70-765.html>2.[2019 Latest 70-765 Exam Questions & Answers Instant Download:

<https://drive.google.com/drive/folders/0B75b5xYLjSSNTnR6dFR2U3A5cFk?usp=sharing>New QuestionYou have a Microsoft Azure SQL database in the US West region.You need to create a replica in the US East region.Which cmdlet should you run first?A. New-AzureRmAvailabilitySetB. New-AzureRmLoadBalancerC. New-AzureRmSqlDatabaseSecondaryD.

New-AzureRmSqlElasticPoolE. New-AzureRmVMF. New-AzureRmSqlServerG. New-AzureRmSqlDatabaseCopyH.

New-AzureRmSqlServerCommunicationLinkAnswer: GExplanation:The New-AzureRmSqlDatabaseCopy command creates a copy of a SQL Database that uses the snapshot at the current time.References:

<https://docs.microsoft.com/en-us/powershell/module/azurermsql/new-azurermsqldatabasecopy?view=azurermps-5.1.1>New

QuestionYou have a database named DB1 that uses simple recovery mode.Full backups of DB1 are taken daily and DB1 is checked for corruption before each backup.There was no corruption when the last backup was complete.You run the sys.columns catalog view and discover corrupt pages.You need to recover the database. The solution must minimize data loss.What should you do?A.

Run RESTORE DATABASE WITH RECOVERY.B. Run RESTORE DATABASE WITH PAGE.C. Run DBCC CHECKDB and specify the REPAIR_ALLOW_DATA_LOSS parameter.D. Run DBCC CHECKDB and specify the REPAIT_REBUILD

parameter.Answer: BExplanation:A page restore is intended for repairing isolated damaged pages. Restoring and recovering a few individual pages might be faster than a file restore, reducing the amount of data that is offline during a restore operation.RESTORE DATABASE WITH PAGERestores individual pages. Page restore is available only under the full and bulk-logged recovery models.

References: **<https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>**New QuestionYou have a database named DB1.You discover that DB1 is corrupt.You run DBCC CHECKDB and receive an error message within a few seconds. No pages are listed in the error message.You need to repair the database corruption as quickly as possible. The solution must minimize data loss.What should you do?A. Run DBCC CHECKDB ('db1', REPAIR_ALLOW_DATA_LOSS).B. Run DBCC CHECKDB ('db1', REPAIR_FAST).C. Delete the transaction logs and restart the Microsoft SQL Server instance.D. Run DBCC CHECKDB ('db1', REPAIR_REBUILD).E. Restore the database from a backup.

Answer: CNew QuestionUser report that a query takes a long time to execute. The query has the following wait statistics.Which resource causes the issue?A. processorB. diskC. blockingD. networkAnswer: BExplanation:PAGEIOLATCH Wait time and WaitCount are both high.One of the most common wait type seen on SQL Server and definitely one that causes a lot of troubles to less experienced database administrators is the PAGEIOLATCH_SH wait type. This is one of those wait types that clearly indicates one thing, but which background and potential causes are much subtler and may lead to erroneous conclusions and worse, incorrect solutions The Microsoft definition of this wait type is:Occurs when a task is waiting on a latch for a buffer that is in an I/O request. The latch request is in Shared mode.

Long waits may indicate problems with the disk subsystem.References: **https://www.sqlshack.com/handling-excessive-sql-server-pageiolatch_sh-wait-types/**New QuestionYou have an on-premises Microsoft SQL server that has a database named DB1. DB1 contains several tables that are stretched to Microsoft Azure.From SQL Server Management Studio (SSMS), a junior database administrator accidentally deletes several rows from the Azure SQL database and breaks the connection to Azure.You need to resume Stretch Database operations.Which two stored procedures should you use? Each correct answer presents part of the solution.NOTE: Each correct selection is worth one point.A. sys.sp_rda_reconcile_batch B. sys.sp_rda_reconcile_indexesC. sys.sp_rda_reauthorize_dbD. sys.sp_rda_reconcile_columnsE. sys.sp_rda_set_rpo_durationAnswer: CDEExplanation:sys.sp_rda_reauthorize_db restores the authenticated connection between a local database enabled for Stretch and the remote database.If you have accidentally deleted columns from the remote table, run sp_rda_reconcile_columns to add columns to the remote table that exist in the Stretch-enabled SQL Server table but not in the remote table.Incorrect Answers:A: sys.sp_rda_reconcile_batch reconciles the batch ID stored in the Stretch-enabled SQL Server table with the batch ID stored in the remote Azure table.Typically you only have to run sp_rda_reconcile_batch if you have manually deleted the most recently migrated data from the remote table. When you manually delete remote data that includes the most recent batch, the batch IDs are out of sync and migration stops.References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-rda-reconcile-batch-transact-sql> **<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-rda-reauthorize-db-transact-sql>**

New QuestionYou plan to deploy Microsoft SQL Server on a Microsoft Azure Virtual machine. The virtual machine will have a

30-TB database and will have 10 1-TB VHDs for the database. You need to configure the storage to meet the following requirements:- Evenly distribute read and write operations across the VHDs.- Minimize the read and write time. Which storage configuration should you use? A. a parity storage pool B. a simple storage pool C. a mirrored storage pool D. a striped volume E. a RAID-5 volume
Answer: D
Explanation: Data that is written to a striped volume is interleaved to all disks at the same time instead of sequentially. Therefore, disk performance is the fastest on a RAID 0 volume as compared to any other type of disk configuration.
Reference:

<https://support.microsoft.com/en-us/help/323433/how-to-establish-a-striped-volume-raid-0-in-windows-server-2003> New

Question You have a database named DB1 that contains a table named Table1. Table1 has 1 billion rows. You import 10 million rows of data into Table1. After the import, users report that queries take longer than usual to execute. You need to identify whether an out-of-date execution plan is causing the performance issue. Which dynamic management view should you use? A.

sys.dm_xtp_transaction_stats B. sys.dm_exec_input_buffer C. sys.dm_db_index_operational_stats D.

sys.dm_db_stats_properties
Answer: D
Explanation:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-db-stats-properties-transaction-sql> New

Question 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. You have Microsoft SQL Server on a Microsoft Azure virtual machine that has a database named DB1. You discover that DB1 experiences WRITE_LOG waits that are longer than 50 ms. You need to reduce the WRITE_LOG wait time. Solution: Move the transaction logs to a faster disk. Does this meet the goal? A. Yes B.

No
Answer: A
Explanation: In SQL Server, if we have a transactional based system and find a high WRITELOG wait type this is a performance bottleneck and can cause the transaction log file to grow rapidly and frequently. It is being recommended to SQL server users that they must archive the log files on a separate disk for getting better performance. References:

<https://atdhebuja.wordpress.com/2016/06/20/resolving-sql-server-transaction-log-waits/> New
Question 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. You have Microsoft SQL Server on a Microsoft Azure virtual machine that has a database named DB1. You discover that DB1 experiences WRITE_LOG waits that are longer than 50 ms. You need to reduce the WRITE_LOG wait time. Solution: Add additional data files to DB1. Does this meet the goal? A. Yes B. No
Answer: B
Explanation: In SQL Server, if we have a transactional based system and find a high WRITELOG wait type this is a performance bottleneck and can cause the transaction log file to grow rapidly and frequently. It is being recommended to SQL server users that they must archive the log files on a separate disk for getting better performance. References:

<https://atdhebuja.wordpress.com/2016/06/20/resolving-sql-server-transaction-log-waits/> New
Question 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. You have Microsoft SQL Server on a Microsoft Azure virtual machine that has a database named DB1. You discover that DB1 experiences WRITE_LOG waits that are longer than 50 ms. You need to reduce the WRITE_LOG wait time. Solution: Add additional log files to tempdb. Does this meet the goal? A. Yes B. No
Answer: B
Explanation: In SQL Server, if we have a transactional based system and find a high WRITELOG wait type this is a performance bottleneck and can cause the transaction log file to grow rapidly and frequently. It is being recommended to SQL server users that they must archive the log files on a separate disk for getting better performance. References:

<https://atdhebuja.wordpress.com/2016/06/20/resolving-sql-server-transaction-log-waits/> !!!RECOMMEND!!!
1. | 2019 Latest 70-765 Exam Dumps (PDF & VCE) Instant Download: <https://www.braindump2go.com/70-765.html> 2. | 2019 Latest 70-765 Study Guide Video Instant Download: YouTube Video: [YouTube.com/watch?v=Qzw2xt-YE10](https://www.youtube.com/watch?v=Qzw2xt-YE10)

2. | 2019 Latest 70-765 Study Guide Video Instant Download: YouTube Video: [YouTube.com/watch?v=Qzw2xt-YE10](https://www.youtube.com/watch?v=Qzw2xt-YE10)