

[May-2018-New] 100% Real Exam Questions-Braindump2go 70-765 PDF and VCE Dumps 155Q Download[119-129]

2018 May New Microsoft 70-765 Exam Dumps with PDF and VCE Just Updated Today! Following are some new 70-765 Real Exam Questions:1.|2018 Latest 70-765 Exam Dumps (PDF & VCE) 155Q

Download:<https://www.braindump2go.com/70-765.html>2.|2018 Latest 70-765 Exam Questions & Answers

Download:<https://drive.google.com/drive/folders/0B75b5xYLjSSNTnR6dFR2U3A5cFk?usp=sharing>QUESTION 119You manage a Microsoft SQL Server instance named SQL1 that has 32 gigabytes (GB) of total memory. The instance supports an app named App1 that only uses a single thread. App1 frequently queries the database using the same index. The operating system and App1 combined require 8 GB of memory to function. You need to ensure that the SQL Server does not limit the performance of App1. What configuration option should you set?A. min memory per query to 4 GBB. index create memory to 16 GBC. max worker threads to 1D. max server memory to 16 GBAnswer: BExplanation: The index creates memory option controls the maximum amount of memory initially allocated for sort operations when creating indexes. The default value for this option is 0 (self-configuring). If more memory is later needed for index creation and the memory is available, the server will use it; thereby, exceeding the setting of this option. If additional memory is not available, the index creation will continue using the memory already allocated. References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-the-index-create-memory-server-configuration-option>QUESTION 120Note: 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 have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in an Always On Availability Group. You expect to have less than 1 million IO transaction per month. You need to recommend a storage solution for the SQL Servers. The solution must minimize costs. Which storage option should you use?A. Premium P10 disk storageB. Premium P20 disk storageC. Premium P30 disk storageD. Standard locally redundant disk storageE. Standard geo-redundant disk storageF. Standard zone redundant blob storageG. Standard locally redundant blob storageH. Standard geo-redundant blob storageAnswer: AExplanation: P10 has 500 IOPS per disk, which provides for more than 1 million IOPS per month. Note: $3600*30*500$ is 54 million IOPS/month.

References: <https://azure.microsoft.com/en-us/pricing/details/managed-disks/>QUESTION 121You create a new Microsoft Azure subscription. You need to create a group of Azure SQL databases that share resources. Which cmdlet should you run first?A. New-AzureRmAvailabilitySetB. New-AzureRmLoadBalancerC. New-AzureRmSqlDatabaseSecondaryD.

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

New-AzureRmSqlServerCommunicationLinkAnswer: DExplanation: SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single Azure SQL Database server and share a set number of resources (elastic Database Transaction Units (eDTUs)) at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database. References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool>QUESTION 122You plan to create an AlwaysOn availability group that will have two replicas in Microsoft Azure and two on-premises replicas. You need to configure the network to support the availability group listener. Which cmdlet should you run first?A. New-AzureRmAvailabilitySetB.

New-AzureRmLoadBalancerC. New-AzureRmSqlDatabaseSecondaryD. New-AzureRmSqlElasticPoolE. New-AzureRmVMF. New-AzureRmSqlServerG. New-AzureRmSqlDatabaseCopyH.

New-AzureRmSqlServerCommunicationLinkAnswer: BExplanation: An availability group listener is a virtual network name that clients connect to for database access. On Azure virtual machines, a load balancer holds the IP address for the listener. The load balancer routes traffic to the instance of SQL Server that is listening on the probe port. Usually, an availability group uses an internal load balancer. References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-portal-sql-ps-alwayson-int-listener>

QUESTION 123You 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/azurerm.sql/new-azurermsqlcopy>

azurermsqldatabasecopy?view=azurermps-5.1.1 QUESTION 124 You 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 REPAIR_REBUILD parameter.

Answer: B
Explanation: 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 PAGE Restores 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> QUESTION 125 You 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: C
Explanation: 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.

Incorrect Answers:
A: The REPAIR_ALLOW_DATA LOSS option is a supported feature but it may not always be the best option for bringing a database to a physically consistent state. If successful, the REPAIR_ALLOW_DATA LOSS option may result in some data loss. In fact, it may result in more data lost than if a user were to restore the database from the last known good backup.
B: REPAIR_FAST maintains syntax for backward compatibility only. No repair actions are performed.
E: Restoring from backup is not the fastest solution.

References:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql> QUESTION 126 You manage an on-premises Microsoft SQL server that has a database named DB1. An application named App1 retrieves customer information for DB1. Users report that App1 takes an unacceptably long time to retrieve customer records. You need to find queries that take longer than 400 ms to run. Which statement should you execute?
A. B. C. D.

Answer: B
Explanation:
Total_worker_time: Total amount of CPU time, reported in microseconds (but only accurate to milliseconds), that was consumed by executions of this plan since it was compiled.

Incorrect Answers:
A: Qs.max_worker_time: Maximum CPU time, reported in microseconds (but only accurate to milliseconds), that this plan has ever consumed during a single execution.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-query-stats-transact-sql> QUESTION 127 User report that a query takes a long time to execute. The query has the following wait statistics. Which resource causes the issue?
A. processor
B. disk
C. blocking
D. network

Answer: B
Explanation: 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 QUESTION 128 You 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_indexes
C. sys.sp_rda_reauthorize_db
D. sys.sp_rda_reconcile_columns

sys.sp_rda_set_rpo_duration
Answer: CD
Explanation: 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>

QUESTION 129 You 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 poolB. a simple storage poolC. a mirrored storage poolD. a striped volumeE. a RAID-5 volumeAnswer: DExplanation: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>

!!!RECOMMEND!!!1.|2018 Latest 70-765 Exam Dumps (PDF & VCE) 155Q

Download:<https://www.braindump2go.com/70-765.html>2.|2018 Latest 70-765 Study Guide Video: YouTube Video:

[YouTube.com/watch?v=_7PKmZXAEfA](https://www.youtube.com/watch?v=_7PKmZXAEfA)