

[March-2019-New70-765 PDF and VCE(Full Version)192Q Download in 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 use Microsoft SQL Server 2012 to develop a database application.You need to create an object that meets the following requirements:- Takes an input variable- Returns a table of values- Cannot be referenced within a viewWhich object should you use?A. Scalar-valued functionB. Inline functionC. User-defined data typeD. Stored procedureAnswer: DExplanation:Stored procedures accept input parameters and return multiple values in the form of output parameters to the calling program. They cannot be used in views.

<https://docs.microsoft.com/en-us/sql/relational-databases/stored-procedures/stored-procedures-database-engine>
New QuestionYou administer a SQL Server 2012 server that contains a database named SalesDb. SalesDb contains a schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales. UserA is granted the Select permission on the Regions table. The Sales role is granted the Select permission on the Customers schema.You need to ensure that the following requirements are met:- The Sales role does not have the Select permission on the Customers schema.- UserA has the Select permission on the Regions table.Which Transact-SQL statement should you use?A. REVOKE SELECT ON Schema::Customers FROM UserAB. DENY SELECT ON Object::Regions FROM UserAC. EXEC sp_addrolemember 'Sales', 'UserA'D. DENY SELECT ON Object::Regions FROM SalesE. REVOKE SELECT ON Object::Regions FROM UserAF. DENY SELECT ON Schema::Customers FROM SalesG. DENY SELECT ON Schema::Customers FROM UserAH. EXEC sp_droprolemember 'Sales', 'UserAI. REVOKE SELECT ON Object::Regions FROM SalesJ. REVOKE SELECT ON Schema::Customers FROM SalesAnswer: JExplanation:Use REVOKE to remove the grant or deny of a permission.

<https://docs.microsoft.com/en-us/sql/t-sql/statements/permissions-grant-deny-revoke-azure-sql-data-warehouse-parallel-data-warehouse>
New QuestionYou develop a Microsoft SQL Server 2012 database that contains a heap named OrdersHistorical.You write the following Transact-SQL query:INSERT INTO OrdersHistoricalSELECT * FROM CompletedOrdersYou need to optimize transaction logging and locking for the statement.Which table hint should you use?A. HOLDLOCKB. ROWLOCKC. XLOCKD. UPDLOCKE. TABLOCKAnswer: EExplanation:When importing data into a heap by using the INSERT INTO SELECT <columns> FROM statement, you can enable optimized logging and locking for the statement by specifying the TABLOCK hint for the target table.

<https://docs.microsoft.com/en-us/sql/t-sql/queries/hints-transact-sql-table>
New QuestionYour database contains a table named Purchases.The table includes a DATETIME column named PurchaseTime that stores the date and time each purchase is made. There is a non- clustered index on the PurchaseTime column. The business team wants a report that displays the total number of purchases made on the current day.You need to write a query that will return the correct results in the most efficient manner.Which Transact-SQL query should you use?A. SELECT COUNT(*)FROM PurchasesWHERE PurchaseTime = CONVERT(DATE, GETDATE())B. SELECT COUNT(*)FROM PurchasesWHERE PurchaseTime = GETDATE()C. SELECT COUNT(*)FROM PurchasesWHERE CONVERT(VARCHAR, PurchaseTime, 112) =CONVERT(VARCHAR, GETDATE(), 112)D. SELECT COUNT(*)FROM PurchasesWHERE PurchaseTime >= CONVERT(DATE, GETDATE())AND PurchaseTime <DATEADD(DAY, 1, CONVERT(DATE, GETDATE()))Answer: DExplanation:To compare a time with date we must use >= and > operators, and not the = operator.Incorrect Answers:A: The in WHERE clause there is an incorrect comparison between time and a date, as equality (=) is used.http://technet.microsoft.com/en-us/library/ms181034.aspxNew QuestionYou have Microsoft SQL Server on a Microsoft Azure virtual machine.You suspect that the current SQL Server indexes cause queries to execute slowly.You need to identify which indexes must be created to reduce the query execution time.Which three dynamic management views should you use? Each correct answer presents part of the solution.NOTE: Each correct selection is worth one point.A. sys.dm_db_index_physical_statsB. sys.dm_db_missing_index_group_statsC. sys.indexesD. sys.dm_db_index_usage_statsE. sys.dm_db_missing_index_groupsF. sys.dm_db_index_operational_statsG. sys.dm_db_missing_index_detailsH. sys.sysindexkeysAnswer: BEExplanation:The missing indexes feature consists of the following components:A set of dynamic management objects that can be queried to return information about missing indexes.The MissingIndexes element in XML Showplans, which correlate indexes that the query optimizer considers missing with the queries for which they are missing.Dynamic Management ObjectsAfter running a typical workload on SQL Server, you can retrieve information about missing indexes by querying the dynamic management objects listed in the following table. These dynamic management objects are stored in the master database.sys.dm_db_missing_index_group_stats

Returns summary information about missing index groups, for example, the performance improvements that could be gained by implementing a specific group of missing indexes.
`sys.dm_db_missing_index_groups`Returns information about a specific group of missing indexes, such as the group identifier and the identifiers of all missing indexes that are contained in that group.
`sys.dm_db_missing_index_details`Returns detailed information about a missing index; for example, it returns the name and identifier of the table where the index is missing, and the columns and column types that should make up the missing index.
`sys.dm_db_missing_index_columns`Returns information about the database table columns that are missing an index.
References: [https://technet.microsoft.com/en-us/library/ms345524\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms345524(v=sql.105).aspx)
New Question You use Microsoft Azure Resource Manager to deploy two new Microsoft SQL Server instances in an Azure virtual machine (VM). VM has 28 gigabytes (GB) of memory. The instances are named Instance1 and Instance2, respectively. The various databases on the instances have the following characteristics: You run the following Transact-SQL statements: You need to configure each SQL Server instance to correctly allocate memory. What should you do?
A. On Instance1, run the following Transact-SQL code: On Instance2, run the following Transact-SQL code:
B. On Instance1, run the following Transact-SQL code: On Instance2, run the following Transact-SQL code:
C. On Instance1, run the following Transact-SQL code: On Instance2, run the following Transact-SQL code:
D. On Instance1, run the following Transact-SQL code: On Instance2, run the following Transact-SQL code:
Answer: D
Explanation: Incorrect Answers:
A, C: The `awe` feature will be removed in newer version of Microsoft SQL Server. Do not use this feature in new development work, and modify applications that currently use this feature as soon as possible.
B: Automatic soft-NUMA is disabled by default, and this setting is not relevant here.
References: <https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/server-configuration-options-sql-server>
New Question You manage a Microsoft SQL Server environment in a Microsoft Azure virtual machine. You must enable Always Encrypted for columns in a database. You need to configure the key store provider. What should you do?
A. Use the Randomized encryption type
B. Modify the connection string for applications
C. Auto-generate a column master key
D. Use the Azure Key Vault
Answer: D
Explanation:
<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/create-and-store-column-master-keys-always-encrypted>
New Question You 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 GB
B. `index create memory` to 16 GB
C. `max worker threads` to 1
D. `max server memory` to 16 GB
Answer: B
Explanation: The `index create 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>
New Question Note: 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 storage
B. Premium P20 disk storage
C. Premium P30 disk storage
D. Standard locally redundant disk storage
E. Standard geo-redundant disk storage
F. Standard zone redundant blob storage
G. Standard locally redundant blob storage
H. Standard geo-redundant blob storage
Answer: A
Explanation: 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/>
New Question You 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-AzureRmAvailabilitySet`
B. `New-AzureRmLoadBalancer`
C. `New-AzureRmSqlDatabaseSecondary`
D. `New-AzureRmSqlElasticPool`
E. `New-AzureRmVMF`
F. `New-AzureRmSqlServer`
G. `New-AzureRmSqlDatabaseCopy`
H. `New-AzureRmSqlServerCommunicationLink`
Answer: D
Explanation: 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> New Question You 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-AzureRmAvailabilitySet B.

New-AzureRmLoadBalancer C. New-AzureRmSqlDatabaseSecondary D. New-AzureRmSqlElasticPool E. New-AzureRmVM F. New-AzureRmSqlServer G. New-AzureRmSqlDatabaseCopy H. New-AzureRmSqlServerCommunicationLink Answer: B Explanation: 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-](https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-portal-sql-ps-alwayson-int-listener)

[windows-portal-sql-ps-alwayson-int-listener](https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-portal-sql-ps-alwayson-int-listener)!!!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-YEI0](https://www.youtube.com/watch?v=Qzw2xt-YEI0)