

## [2018-June-NewHigh Quality Braindump2go 70-764 Exam Dumps VCE and PDF 365Q Free Share[220-230

2018 June New Microsoft 70-764 Exam Dumps with PDF and VCE Free Updated Today! Following are some new 70-764 Real Exam Questions:1.|2018 New 70-764 Exam Dumps (PDF & VCE) 365Q&As

Download:<https://www.braindump2go.com/70-764.html>2.|2018 New 70-764 Exam Questions & Answers

Download:<https://drive.google.com/drive/folders/0B75b5xYLjSSNdlF6dzFQVE9kUjA?usp=sharing> QUESTION 220A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases. A customer that uses an on-premises instance reports that queries take a long time to complete. You need to reconfigure table statistics so that the query optimizer can use the optimal query execution plans available. Which Transact-SQL segment should you use?A. sys.index\_columnsB. UPDATE STATISTICSC. CREATE STATISTICSD. SET AUTO\_CREATE\_STATISTICS ONAnswer: DExplanation: AUTO\_UPDATE\_STATISTICS { ON | OFF }ON specifies that the query optimizer updates statistics when they are used by a query and when they might be out-of-date. Statistics become out-of-date after insert, update, delete, or merge operations change the data distribution in the table or indexed view. The query optimizer determines when statistics might be out-of-date by counting the number of data modifications since the last statistics update and comparing the number of modifications to a threshold. The threshold is based on the number of rows in the table or indexed view. References: [https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-transact-sql-set-options?view=sql-server-2017#auto\\_update\\_statistics](https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-transact-sql-set-options?view=sql-server-2017#auto_update_statistics)

QUESTION 221Hotspot QuestionA company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environments host several customer databases. You configure an Always On Availability Group for a customer. You must create log reports for the customer that detail when the log is flushed to disk on the primary and secondary replica. You need to develop a report containing the requested information. In the table below, identify the log type that you should use for each replica. NOTE: Make only one selection in each column. Each correct selection is worth one point. Answer: Explanation: Flush on primary: Log flushLog flush. Log data is generated and flushed to disk on the primary replica in preparation for replication to the secondary replica. It then enters the send queue. Flush on secondary: Log hardenedThe log is flushed on the secondary replica, and then a notification is sent to the primary replica to acknowledge completion of the transaction. Incorrect Answers: Not Log capture: Log capture. Logs for each database are captured on the primary replica, compressed, and sent to the corresponding queue on the secondary replica. This process runs continuously as long as database replicas are connecting. If this process is not able to scan and enqueue the messages quickly enough, the log send queue continues to grow. Not Log receive/Log cacheLog receive/Log cache. Each secondary replica gets messages from the primary replica and then caches the messages. References:

[http://www.futas.net/ora/doc/SQL\\_Server\\_2016\\_Higher\\_availability\\_eBook\\_EN\\_US.pdf](http://www.futas.net/ora/doc/SQL_Server_2016_Higher_availability_eBook_EN_US.pdf) QUESTION 222Drag and Drop QuestionYou are designing a high availability (HA) environment for a company that has three office locations. Details of the services deployed at each office are shown in the table below: You need to maximize availability, minimize data loss, and minimize downtime in the event of a failure. Which solution should you implement for each location? To answer, drag the appropriate solutions to the correct locations. Each solution 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. Answer: Explanation: The Always On availability groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. References:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/always-on-availability-groups-sql-server?view=sql-server-2017> QUESTION 223Drag and Drop QuestionYou have a database named DB1. You must ensure that DB1 is highly available. You implement a Windows cluster. You need to create a Cluster Shared Volume (CSV) for the database and log files. Which three Windows Power Shell commands should you execute in sequence? To answer, move the appropriate Windows PowerShell commands from the list of Windows PowerShell commands to the answer area and arrange them in the correct order. Answer: Explanation: Box 1: \$cluster = Get-ClusterThe Get-Cluster cmdlet gets information about one or more failover clusters in a given domain. This cmdlet can obtain a variety of configuration and state information about a failover cluster, including the following items, such as the EnabledSharedVolumes property. Box 2: \$cluster.EnabledSharedVolumes ="Enabled" Box 3: Add-ClusterSharedVolume "Disk1" The Add-ClusterSharedVolume cmdlet makes a volume available in the Cluster Shared Volumes in a failover cluster. The Cluster Shared Volume must be chosen from the disks in Available Storage (that is, disks that have been added to the cluster but not assigned to a specific use in the cluster). References:

<https://docs.microsoft.com/en-us/powershell/module/failoverclusters/get-cluster>

<https://docs.microsoft.com/en-us/powershell/module/failoverclusters/add-clustersharedvolume?view=winserver2012-ps>

**QUESTION 224** Hotspot Question You are planning the deployment of two new Always On Failover Cluster Instances (FCIs) of Microsoft SQL Server to a single Windows Server Cluster with three nodes. The planned configuration for the cluster is shown in the Server Layout exhibit. (Click the Exhibit button.) The SAN team has configured storage for the cluster and sent the configuration to you in the email shown in the SAN Team Email exhibit. (Click the Exhibit button.) Each node of the cluster has identical local storage available as shown in the Local Storage exhibit. (Click the Exhibit button.) All local storage is on SSD. You need to plan specific configurations for the new cluster. For each of the following statement, select Yes if the statement is true. Otherwise, select No. Answer: Explanation: Box 1: Yes tempdb on local storage. FCIs now support placement of tempdb on local non-shared storage, such as a local solid-state-drive, potentially offloading a significant amount of I/O from a shared SAN. Prior to SQL Server 2012, FCIs required tempdb to be located on a symmetrical shared storage volume that failed over with other system databases. Box 2: No The VNN is set on the group level, not on the instance level. Database client applications can connect directly to a SQL Server instance network name, or they may connect to a virtual network name (VNN) that is bound to an availability group listener. The VNN abstracts the WSFC cluster and availability group topology, logically redirecting connection requests to the appropriate SQL Server instance and database replica. The logical topology of a representative AlwaysOn solution is illustrated in this diagram:

**QUESTION 225** Hotspot Question 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. Start of repeated scenario. You are a database administrator for a company that has on-premises Microsoft SQL Server environment. There are two domains in separate forests. There are no trust relationships between the domains. The environment hosts several customer databases, and each customer uses a dedicated instance running SQL Server 2016 Standard edition. The customer environments are shown in the following table. End of repeated scenario. You need to configure auditing for the AdventureWorks environment. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point. Answer: Explanation: Scenario: You must implement auditing for all objects in the ADVSchema schema. Box 1: CREATE SERVER AUDIT Create the server audit. Box 2: ALTER SERVER AUDIT Enable the server audit. Box 3: CREATE DATABASE AUDIT Create the database audit specification. Box 4: FOR SERVER AUDIT Example: The following example creates a server audit called Payrole\_Security\_Audit and then a database audit specification called Payrole\_Security\_Audit that audits SELECT and INSERT statements by the dbo user, for the HumanResources.EmployeePayHistory table in the AdventureWorks2012 database. USE master ; GO -- Create the server audit.

```
CREATE SERVER AUDIT Payrole_Security_Audit TO FILE ( FILEPATH = 'C:Program FilesMicrosoft SQL ServerMSSQL13.MSSQLSERVERMSSQLDATA' ) ; GO -- Enable the server audit. ALTER SERVER AUDIT Payrole_Security_Audit WITH ( STATE = ON ) ; GO -- Move to the target database. USE AdventureWorks2012 ; GO -- Create the database audit specification. CREATE DATABASE AUDIT SPECIFICATION Audit_Pay_Tables FOR SERVER AUDIT Payrole_Security_Audit ADD ( SELECT , INSERT ON HumanResources.EmployeePayHistory BY dbo ) WITH ( STATE = ON ) ; GO References:
```

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-audit-specification-transact-sql?view=sql-server-2017>

**QUESTION 226** Hotspot Question 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. Start of repeated scenario. You are a database administrator for a company that has on-premises Microsoft SQL Server environment. There are two domains in separate forests. There are no trust relationships between the domains. The environment hosts several customer databases, and each customer uses a dedicated instance running SQL Server 2016 Standard edition. The customer environments are shown in the following table. End of repeated scenario. You need to configure auditing for WDWDB. In the table below, identify the event type that you must audit for each activity. NOTE: Make only one selection in each column. Answer: Explanation: You must log any deletion of views and all database record update operations.

**QUESTION 227** Hotspot Question You manage a Microsoft SQL Server environment with a database named salesOrders. You plan to encrypt data during database backups. You need to ensure that you can recover the database encryption keys. How should you complete the Transact-SQL statement? To answer, select the appropriate Transact-SQL segments in the answer area. Answer: Explanation: -- Creates a backup of the "AdventureWorks2012" master key. Because this master key is not encrypted by the service master key, a password must be specified when it is opened. USE AdventureWorks2012; GO OPEN MASTER KEY DECRYPTION BY PASSWORD = 'sfj5300osdVdgwdfkli7'; BACKUP MASTER KEY TO FILE = 'c:tempexportedmasterkey' ENCRYPTION BY PASSWORD = 'sd092735kjn\$&adsg'; GO References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/back-up-a-database-master-key?view=sql-server-2017>

**QUESTION 228** Drag and Drop Question You are deploying a Microsoft Azure SQL Database environment. You need to

configure the environment to meet the following requirements: Access to databases must be logged. Email alerts must be sent to the database administrator if anomalous activity is detected. What should you do? To answer, drag the appropriate solutions to the correct requirements. Each solution 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. Answer: Explanation:Box 1: Enable database auditingBox 2: Enable threat detectionSQL Database Threat Detection runs multiple sets of algorithms which detect potential vulnerabilities and SQL injection attacks, as well as anomalous database access patterns (such as access from an unusual location or by an unfamiliar principal). Security officers or other designated administrators get email notification once a threat is detected on the database. Each notification provides details of the suspicious activity and recommends how to further investigate and mitigate the threat. Reference:

[https://azure.microsoft.com/en-us/blog/azure-sql-database-threat-detection-general-availability-in-spring-2017/QUESTION\\_229](https://azure.microsoft.com/en-us/blog/azure-sql-database-threat-detection-general-availability-in-spring-2017/QUESTION_229)  
Drag and Drop Question You have a database named DB1 that contains the primary file group and two other file groups named CurrentData and HistoricalData. CurrentData contains 320 GB of active data that is updated frequently. HistoricalData is read-only and contains 1.7 TB of data. You have a backup schedule configured as shown in the following table. DB1 fails at 08:42 on Monday. You deploy a new server. You need to recover CurrentData as quickly as possible on the new server. The solution must minimize the amount of data loss. Which four statements should you execute in sequence? To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order. Answer: Explanation:Step 1: ..FILEGROUP = N'PRIMARY.Restore the primary filegroupStep 2: ..CurrentDataRestore the CurrentDATA filegroup (not the HistoricalDATA).Step 3: RESTORE ..\diff\DB1\_diff..Restore the differential backupStep 4: RESTORE LOGFinally restore the log files.QUESTION 230 Hotspot Question Your company hosts Microsoft SQL Server instances for customers. One customer requires that the SQL Server instance send an email when an alert is generated in the Sales database. You need to create the alert. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area. Answer: Explanation: The following example adds an e-mail notification for the specified alert (Test Alert).name. USE msdb ;GOEXEC dbo.sp\_add\_notification@alert\_name = N'Test Alert',@notification\_method = 1 ;GOREferences:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-notification-transact-sql?view=sql-server-2017>!!!RECOMMENDED!!!]1.|2018 New 70-764 Exam Dumps (PDF & VCE) 365Q&As

Download:<https://www.braindump2go.com/70-764.html>2.|2018 New 70-764 Study Guide Video: YouTube Video:

[YouTube.com/watch?v=1Udu9KJu8](https://www.youtube.com/watch?v=1Udu9KJu8)