

[2018-June-NewHigh Quality Braindump2go 70-764 VCE Exam Dumps 365Q Free Share[209-219

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/0B75b5xYLjSSNdIF6dzFQVE9kUjA?usp=sharing> QUESTION 209Note: 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 are deploying a Microsoft SQL Server architecture to support a new mission-critical application. The application includes a dedicated reporting component that performs read-only operations against the application's databases.You need to implement a solution that meets the following requirements: Provide maximum uptime for the databases.

Include automatic failover in the event of a hardware problem on the primary server. Separate the reporting workload from the read/write transactional processing workload and contain real-time data.Modifications to the application to support the new architecture are not permitted.What should you implement?A. a Microsoft Azure Stretch DatabaseB. log shippingC. an Always On Availability Group with all replicas in synchronous-commit modeD. a file share witnessE. a SQL Server failover cluster instance (FCI)F. a Windows Cluster with a shared-nothing architectureG. an Always On Availability group with secondary replicas in asynchronous-commit modeAnswer: CExplanation:Synchronous-commit mode emphasizes high availability over performance, at the cost of increased transaction latency.References:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on-availability-groups?view=sql-server-2017>QUESTION 210Note: 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 monitor WingDB and gather information for troubleshooting issues.Which two tools should you use? Each correct answer presents a complete solution.A. sys.dm_tran_locksB. sp_lockC. sys.dm_tran_active_snapshot_database_transactionsD. Activity MonitorE. sp_monitorAnswer: BDEExplanation:The performance issues is related to locking.B: sp_lock reports snapshot information about locks, including the object ID, index ID, type of lock, and type or resource to which the lock applies.D: The Activity Monitor in SQL Server Management Studio is useful for ad hoc views of current activity and graphically displays information about:Processes running on an instance of SQL Server.Blocked processes.Locks.User activity.Incorrect Answers:E: System Monitor primarily tracks resource usage, such as the number of buffer manager page requests in use, enabling you to monitor server performance and activity using predefined objects and counters or user-defined counters to monitor events. System Monitor (Performance Monitor in Microsoft Windows NT 4.0) collects counts and rates rather than data about the events (for example, memory usage, number of active transactions, number of blocked locks, or CPU activity). You can set thresholds on specific counters to generate alerts that notify operatorReferences:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/performance-monitoring-and-tuning-tools?view=sql-server-2017>QUESTION 211You have a database named DB1 that contains two tables.You need to encrypt one column in each table by using the Always Encrypted feature.The solution must support groupings on encrypted columns.Which two actions should you perform? Each correct answer presents part of the solution.NOTE: Each correct selection is worth one point.A. Encrypt both columns by using deterministic encryption.B. Provision a symmetric key by using Transact-SQL.C. Encrypt both columns by using randomized encryption.D. Provision column master keys and column encryption keys by using Microsoft SQL Server Management Studio (SSMS).Answer: ADEExplanation:A: Use deterministic encryption for columns that will be used as search or grouping parameters, for example a government ID number.Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns.D: Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.Incorrect Answers:B: A column encryption key (CEK), is a content encryption key (i.e. a key used to protect data) that is protected by a CMK.All Microsoft CMK store providers encrypt CEKs by using RSA with Optimal Asymmetric Encryption

Padding (RSA-OAEP) with the default parameters specified by RFC 8017 in Section A.2.1.C: Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine?view=sql-server-2017> QUESTION 212 You manage a Microsoft SQL Server environment. You plan to encrypt data when you create backups. You need to configure the encryption options for backups. What should you configure? A. a certificate B. an MD5 hash C. an SHA-256 hash D. an AES 256-bit key Answer: D Explanation: To encrypt a backup we need to configure an encryption algorithm (supported encryption algorithms are: AES 128, AES 192, AES 256, and Triple DES) and an encryptor (a certificate or asymmetric key). Example: in order to create an encrypted backup we must create a Database Master Key. USE SampleDBGOCREATE MASTER KEY ENCRYPTION BY PASSWORD = 'MSSQLTips.com'; GO We must create a certificate to act as an encryptor. Use SampleDBGOCREATE CERTIFICATE SampleDB_Backup_Certificate WITH SUBJECT = 'SQL Server 2014'; GO Now we are ready to create encrypted backups. References:

<https://www.mssqltips.com/sqlservertip/3145/sql-server-2014-backup-encryption/> QUESTION 213 You have a database that stores information for a shipping company. You create a table named Customers by running the following Transact-SQL statement. (Line numbers are included for reference only.) The salesPerson column stores the username of the salesperson. You must create a security policy that ensures that salespeople can view data only for the customers that are assigned to them. You need to create the function that will be used as the filter predicate for the security policy. You write the following Transact-SQL: Which Transact-SQL segment should you insert at line 02? A. B. C. D. Answer: B Explanation: The inline table-valued function required for a Stretch Database filter predicate looks like the following example. The parameters for the function have to be identifiers for columns from the table. Schema binding is required to prevent columns that are used by the filter function from being dropped or altered. Example: CREATE FUNCTION dbo.fn_stretchpredicate(@column1 datatype1, @column2 datatype2 [, ...n]) RETURNS TABLE WITH SCHEMABINDING AS RETURN SELECT 1 AS is_eligible WHERE <predicate> References:

<https://docs.microsoft.com/en-us/sql/sql-server/stretch-database/select-rows-to-migrate-by-using-a-filter-function-stretch-database?view=sql-server-2017> QUESTION 214 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. A company has an on-premises Microsoft SQL Server environment. SQL Server backups should be stored as Microsoft Azure blob pages. The connection process from the SQL Server instances to Azure should be encrypted. You need to store backups as Azure blob pages. Which option should you use? A. backup compression B. backup encryption C. file snapshot backup D. mirrored backup media sets E. SQL Server backup to URL F. SQL Server Managed Backup to Azure G. tail-log backup H. back up and truncate the transaction log Answer: F Explanation: SQL Server Managed Backup to Microsoft Azure manages and automates SQL Server backups to Microsoft Azure Blob storage. You can choose to allow SQL Server to determine the backup schedule based on the transaction workload of your database. Or you can use advanced options to define a schedule. The retention settings determine how long the backups are stored in Azure Blob storage. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-managed-backup-to-microsoft-azure?view=sql-server-2017> QUESTION 215 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. A company has several Microsoft SQL Server databases in Microsoft Azure. One database experiences a storage failure, and pages that store critical database metadata are corrupted. You need to perform an offline restore of the database's pages. Which option should you use first? A. backup compression B. backup encryption C. file snapshot backup D. mirrored backup media sets E. SQL Server backup to URL F. SQL Server Managed Backup to Azure G. tail-log backup H. back up and truncate the transaction log Answer: G Explanation: An unbroken chain of log backups must be available, up to the current log file, and they must all be applied to bring the page up to date with the current log file. A tail-log backup captures any log records that have not yet been backed up (the tail of the log) to prevent work loss and to keep the log chain intact. Before you can recover a SQL Server database to its latest point in time, you must back up the tail of its transaction log. The tail-log backup will be the last backup of interest in the recovery plan for the database. Note: The goal of a page restore is to restore one or more damaged pages without restoring the whole database. Typically, pages that are candidates for restore have been marked as "suspect" because of an error that is encountered when accessing the page. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/tail-log-backups-sql-server?view=sql-server-2017#TailLogScenarios>

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/restore-pages-sql-server?view=sql-server-2017#Restrictions>

QUESTION 216Note: 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.A company has a Microsoft SQL Server environment that has multiple databases. A database named DB1 has multiple online file groups. It is configured to use the full recovery model. A full backup is preformed nightly and transaction logs are performed on the hour. A large number of records are accidentally deleted at 17:20.You need to perform a point-in-time recovery. Which option should you use first?A. backup compressionB. backup encryptionC. file snapshot backupD. mirrored backup media setsE. SQL Server backup to URLF. SQL Server Managed Backup to AzureG. tail-log backupH. back up and truncate the transaction logAnswer: GExplanation:To back up the tail of the log (that is, the active log), check Back up the tail of the log, and leave database in the restoring state.A tail-log backup is taken after a failure to back up the tail of the log in order to prevent work loss. Back up the active log (a tail-log backup) both after a failure, before beginning to restore the database, or when failing over to a secondary database. Selecting this option is equivalent to specifying the NORECOVERY option in the BACKUP LOG statement of Transact-SQL.References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/back-up-a-transaction-log-sql-server?view=sql-server-2017>

QUESTION 217Note: 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.A company has a Microsoft SQL Server environment in Microsoft Azure. The databases are stored directly in Azure blob storage.You need to ensure that you can restore a database to a specific point in time between backups while minimizing the number of Azure storage containers required.Which option should you use?A. backup compressionB. backup encryptionC. file snapshot backupD. mirrored backup media setsE. SQL Server backup to URLF. SQL Server Managed Backup to AzureG. tail-log backupH. back up and truncate the transaction logAnswer: FExplanation:SQL Server Managed Backup to Microsoft Azure supports point in time restore for the retention time period specified. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-managed-backup-to-microsoft-azure?view=sql-server-2017>

QUESTION 218Note: 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.A company has a Microsoft SQL Server environment in Microsoft Azure. The databases are stored directly in Azure blob storage. The company uses a complex backup process.You need to simplify the backup process. Future restores should not require differential or multiple incremental logs to perform a restore.You need to design a backup solution for the SQL Server instances.Which option should you use?A. backup compressionB. backup encryptionC. file snapshot backupD. mirrored backup media setsE. SQL Server backup to URLF. SQL Server Managed Backup to AzureG. tail-log backupH. back up and truncate the transaction logAnswer: CExplanation:SQL Server File-snapshot backup uses Azure snapshots to provide nearly instantaneous backups and quicker restores for database files stored using the Azure Blob storage service. This capability enables you to simplify your backup and restore policies.References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/file-snapshot-backups-for-database-files-in-azure?view=sql-server-2017>

QUESTION 219You have an application that queries a database. Users report that the application is slower than expected.You discover that several server process identifiers (SPIDs) have PAGELATCH_UP and PAGELATCH_EX waits. The resource descriptions of the SPIDs contains 2:1:1.You need to resolve the issue.What should you do?A. Allocate additional processor cores to the server.B. Add files to the file group of the application database.C. Reduce the fill factor of all clustered indexes.D. Add data files to tempdb.Answer: DExplanation:PAGELATCH contention in tempdb is typically on allocation bitmaps and occurs with workloads with many concurrent connections creating and dropping small temporary tables (which are stored in tempdb).Assuming that the temporary tables are needed for performance, the trick is to have multiple data files for tempdb so that the allocations are done round-robin among the files, the contention is split over multiple PFS pages, and so the overall contention goes down.References: <https://sqlperformance.com/2015/10/sql-performance/knee-jerk-wait-statistics-pagelatch>

!!!RECOMMEBD!!!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=1Udu9KJuu8](https://www.youtube.com/watch?v=1Udu9KJuu8)