

[2018-June-NewValid Braindump2go 70-764 Exam VCE and PDF 365Q Offer[198-208

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 198Note: 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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You support an application that stores data in a Microsoft SQL Server database. You have a query that returns data for a report that users run frequently.The query optimizer sometimes generates a poorly-performing plan for the query when certain parameters are used. You observe that this is due to the distribution of data within a specific table that the query uses.You need to ensure that the query optimizer always uses the query plan that you prefer.Solution: You force the desired plan.Does the solution meet the goal?A. YesB. NoAnswer: BExplanation:KEEPFIXED PLAN should be used as it forces the query optimizer not to recompile a query due to changes in statistics.When FORCEPLAN is set to ON, the SQL Server query optimizer processes a join in the same order as the tables appear in the FROM clause of a query. In addition, setting FORCEPLAN to ON forces the use of a nested loop join unless other types of joins are required to construct a plan for the query, or they are requested with join hints or query hints.References: <https://docs.microsoft.com/en-us/sql/t-sql/queries/hints-transact-sql-query?view=sql-server>-QUESTION 199Note: 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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You support an application that stores data in a Microsoft SQL Server database. You have a query that returns data for a report that users run frequently.The query optimizer sometimes generates a poorly-performing plan for the query when certain parameters are used. You observe that this is due to the distribution of data within a specific table that the query uses.You need to ensure that the query optimizer always uses the query plan that you prefer.Solution: You create a copy of the plan guide for the query plan.Does the solution meet the goal?A. YesB. NoAnswer: BExplanation: PLAN should be used as it forces the query optimizer not to recompile a query due to changes in statistics.References: <https://docs.microsoft.com/en-us/sql/t-sql/queries/hints-transact-sql-query?view=sql-server>-QUESTION 200Note: 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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You support an application that stores data in a Microsoft SQL Server database. You have a query that returns data for a report that users run frequently.The query optimizer sometimes generates a poorly-performing plan for the query when certain parameters are used. You observe that this is due to the distribution of data within a specific table that the query uses.You need to ensure that the query optimizer always uses the query plan that you prefer.Solution: You add the KEEPFIXED PLAN query hint to the query.Does the solution meet the goal?A. YesB. NoAnswer: AExplanation:PLAN forces the query optimizer not to recompile a query due to changes in statistics. Specifying KEEPFIXED PLAN makes sure that a query will be recompiled only if the schema of the underlying tables is changed or if sp_recompile is executed against those tables.References: <https://docs.microsoft.com/en-us/sql/t-sql/queries/hints-transact-sql-query?view=sql-server>-QUESTION 201Note: 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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You have a data warehouse that stores sales data. One fact table has 100 million rows.You must reduce storage needs for the data warehouse.You need to implement a solution that uses column-based storage and provides real-time analytics for the operational workload.Solution: You remove any clustered indexes and load the table for processing.Does the solution meet the goal?A. YesB. NoAnswer: BExplanation:Clustered columnstore tables offer both the highest level of data compression as well as the best overall query performance. Clustered columnstore tables will generally outperform clustered index or heap tables and are usually the best choice for large tables. For these reasons, clustered columnstore is the best place to start when you are unsure of how to index your table.Note: Dimensional tables can be used to reduce the size of fact tables.Dimension tables contain attribute

data that might change but usually changes infrequently. For example, a customer's name and address are stored in a dimension table and updated only when the customer's profile changes. To minimize the size of a large fact table, the customer's name and address do not need to be in every row of a fact table. Instead, the fact table and the dimension table can share a customer ID. A query can join the two tables to associate a customer's profile and transactions. References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-overview> QUESTION 202 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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. Your company is developing a new business intelligence application that will access data in a Microsoft Azure SQL Database instance. All objects in the instance have the same owner. A new security principal named BI_User requires permission to run stored procedures in the database. The stored procedures read from and write to tables in the database. None of the stored procedures perform IDENTIFY_INSERT operations or dynamic SQL commands. The scope of permissions and authentication of BI_User should be limited to the database. When granting permissions, you should use the principle of least privilege. You need to create the required security principals and grant the appropriate permissions. Solution: You run the following Transact-SQL statement in the master database: CREATE LOGIN BI_User WITH PASSWORD = `Pa\$\$w You run the following Transact-SQL statement in the business intelligence database: Does the solution meet the goal? A. Yes B. No Answer: B Explanation: One method of creating multiple lines of defense around your database is to implement all data access using stored procedures or user-defined functions. You revoke or deny all permissions to underlying objects, such as tables, and grant EXECUTE permissions on stored procedures. This effectively creates a security perimeter around your data and database objects. Best Practices Simply writing stored procedures isn't enough to adequately secure your application. You should also consider the following potential security holes. Grant EXECUTE permissions on the stored procedures for database roles you want to be able to access the data. Revoke or deny all permissions to the underlying tables for all roles and users in the database, including the public role. All users inherit permissions from public. Therefore denying permissions to public means that only owners and sysadmin members have access; all other users will be unable to inherit permissions from membership in other roles. Do not add users or roles to the sysadmin or db_owner roles. System administrators and database owners can access all database objects. References:

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/managing-permissions-with-stored-procedures-in-sql-server> QUESTION 203 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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. Your company is developing a new business intelligence application that will access data in a Microsoft Azure SQL Database instance. All objects in the instance have the same owner. A new security principal named BI_User requires permission to run stored procedures in the database. The stored procedures read from and write to tables in the database. None of the stored procedures perform IDENTIFY_INSERT operations or dynamic SQL commands. The scope of permissions and authentication of BI_User should be limited to the database. When granting permissions, you should use the principle of least privilege. You need to create the required security principals and grant the appropriate permissions. Solution: You run the following Transact-SQL statement in the database: Does the solution meet the goal? A. Yes B. No Answer: B Explanation: One method of creating multiple lines of defense around your database is to implement all data access using stored procedures or user-defined functions. You revoke or deny all permissions to underlying objects, such as tables, and grant EXECUTE permissions on stored procedures. This effectively creates a security perimeter around your data and database objects. Best Practices Simply writing stored procedures isn't enough to adequately secure your application. You should also consider the following potential security holes. Grant EXECUTE permissions on the stored procedures for database roles you want to be able to access the data. Revoke or deny all permissions to the underlying tables for all roles and users in the database, including the public role. All users inherit permissions from public. Therefore denying permissions to public means that only owners and sysadmin members have access; all other users will be unable to inherit permissions from membership in other roles. Do not add users or roles to the sysadmin or db_owner roles. System administrators and database owners can access all database objects. References:

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/managing-permissions-with-stored-procedures-in-sql-server> QUESTION 204 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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. Your company is developing a new business intelligence application that will

access data in a Microsoft Azure SQL Database instance. All objects in the instance have the same owner. A new security principal named BI_User requires permission to run stored procedures in the database. The stored procedures read from and write to tables in the database. None of the stored procedures perform IDENTIFY_INSERT operations or dynamic SQL commands. The scope of permissions and authentication of BI_User should be limited to the database. When granting permissions, you should use the principle of least privilege. You need to create the required security principals and grant the appropriate permissions. Solution: You run the following Transact-SQL statement: Does the solution meet the goal? A. Yes B. No Answer: A Explanation: One method of creating multiple lines of defense around your database is to implement all data access using stored procedures or user-defined functions. You revoke or deny all permissions to underlying objects, such as tables, and grant EXECUTE permissions on stored procedures. This effectively creates a security perimeter around your data and database objects. Best Practices Simply writing stored procedures isn't enough to adequately secure your application. You should also consider the following potential security holes. Grant EXECUTE permissions on the stored procedures for database roles you want to be able to access the data. Revoke or deny all permissions to the underlying tables for all roles and users in the database, including the public role. All users inherit permissions from public. Therefore denying permissions to public means that only owners and sysadmin members have access; all other users will be unable to inherit permissions from membership in other roles. Do not add users or roles to the sysadmin or db_owner roles. System administrators and database owners can access all database objects. References:

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/managing-permissions-with-stored-procedures-in-sql-server> QUESTION 205

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 attempt to restore a database on a new SQL Server instance and receive the following error message: "Msg 33111, Level 16, State 3, Line 2 Cannot find server certificate with thumbprint `0x7315277C70764B1F252DC7A5101F6F66EFB1069D`." You need to ensure that you can restore the database successfully.

Solution: You disable BitLocker Drive Encryption (BitLocker) on the drive that contains the database backup. Does this meet the goal? A. Yes B. No Answer: B Explanation: This is a certificate problem. The problem is not related to BitLocker. References:

<https://www.sqlservercentral.com/Forums/Topic1609923-3411-1.aspx> QUESTION 206

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 have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram. You have an Always On Availability group named AG1. The details for AG1 are shown in the following table. Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size. Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O. Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes. You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain. You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1. All databases use the full recovery model. All backups are written to the network location \SQLBackup. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table. Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION. You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1. Reporting system: This solution accesses data in DB1 with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas. Operations system: This solution accesses data in DB1 with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations. The wait statistics monitoring requirements for the instances are described in the following table. End of repeated scenario. You need to create a backup plan for Instance4. Which backup plan should you create? A.

Weekly full backups, nightly differential backups, transaction log backups every 30 minutes. B. Weekly full backups, nightly differential. No transaction log backups are necessary. C. Weekly full backups, nightly differential backups, transaction log

backups every 12 hours.D. Full backups every 60 minutes, transaction log backups every 30 minutes.Answer: AExplanation: Scenario: Instance4 is engaged in heavy read-write I/O. The Recovery Point Objective of Instance4 is 60 minutes.QUESTION 207 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 have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram. You have an Always On Availability group named AG1. The details for AG1 are shown in the following table. Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.All databases use the full recovery model. All backups are written to the network location \SQLBackup. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table. Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1. Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas. Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.The wait statistics monitoring requirements for the instances are described in the following table. End of repeated scenario. You need to reduce the amount of time it takes to back up OperationsMain.What should you do? A. Modify the backup script to use the keyword NO_COMPRESSION in the WITH statement.B. Modify the backup script to use the keywords INIT and SKIP in the WITH statement.C. Run the following Transact-SQL statement for each file in OperationsMain: BACKUP DATABASE OperationsMain FILE [...]D. Run the following Transact-SQL statement: BACKUP DATABASE OperationsMain READ_WRITE_FILEGROUPSAnswer: DExplanation: READ_WRITE_FILEGROUPS specifies that all read/write filegroups be backed up in the partial backup. If the database is read-only, READ_WRITE_FILEGROUPS includes only the primary filegroup.Scenario: Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.Incorrect Answers:A: Because a compressed backup is smaller than an uncompressed backup of the same data, compressing a backup typically requires less device I/O and therefore usually increases backup speed significantly.B: INIT and SKIP would not affect backup speed.References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/backup-transact-sql?view=sql-server>QUESTION 208Note: 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 migrating a set of databases from an existing Microsoft SQL Server instance to a new instance. You need to complete the migration while minimizing administrative effort and downtime.Which should you implement?A. log shippingB. an Always On Availability Group with all replicas in synchronous-commit modeC. a file share witnessD. a SQL Server failover cluster instance (FCI)E. a Windows Cluster with a shared-nothing architectureF. an Always On Availability Group with secondary replicas in asynchronous-commit modeAnswer: AExplanation:SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually.References: <https://docs.microsoft.com/en-us/sql/database-engine/log-shipping/about-log-shipping-sql-server?view=sql-server-2017> !!!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)