

[2018-March-NewValid Braindump2go 70-764 PDF and VCE Dumps 332Q Offer[183-193

2018 March New Microsoft 70-764 Exam Dumps with PDF and VCE Free Updated Today! Following are some new 70-764 Real Exam Questions:1.|2018 Latest 70-764 Exam Dumps (PDF & VCE) 332Q&As Download:

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

<https://drive.google.com/drive/folders/0B75b5xYLjSSndIF6dzFQVE9kUjA?usp=sharing> QUESTION 183Your company stored customer data, including credit card numbers, in a Microsoft SQL Server database. The CreditCardNum column is configured as a varchar(16). When viewing the CreditCardNum column, only the last four digits of the card number should be displayed.You have the following Transact-SQL statement. (Line numbers are included for reference only.) You need to implement dynamic data masking for the CreditCardNum column.Which Transact-SQL segment should you insert at line 05?A. B. C. D. Answer: C Explanation:Example:Using "partial" functionALTER TABLE Ari_Users_Tbl ALTER COLUMN [CreditCard]ADD MASKED WITH (FUNCTION = 'partial(0,"XXXX-XXXX-XXXX-",4)') References:

https://social.technet.microsoft.com/wiki/contents/articles/35003.sql-server-exposing-masked-data.aspx#Goal_2_Find_the_CreditCard_number_for_UserID 1QUESTION 184Hotspot QuestionNote: 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.You maintain a Microsoft SQL Server instance that contains the following databases SalesDb1, SalesDb2, and SalesDb3. Each database has table named Products and Sales. The following table shows the configuration of each database. The backup strategies for each database are described in the following table. Each full or differential backup operation writes into a new file and uses a different sequence number. You observe the following database corruption issues. SalesDb3 reports a number of database corruption issues related to error 823 and 824 when reading data pages. You must display the following information about the corrupted pages:- database name- impacted file id- impacted file physical name- impacted page id- event type that identifies the error type- error countUsers report performance issues when they run queries against SalesDb2. You plan to monitor query statistics and execution plans for SalesDb2 by using Query Store. The monitoring strategy must meet the following requirements:- Perform automatic data cleanup when query store disk usage reaches 500 megabyte (MB).- Capture queries based on resource consumption.- Use a stale query threshold value of 60 days.The query optimizer generates suboptimal execution plans for a number of queries on the Sales table in SalesDb2. You will create a maintenance plan that updates statistics for the table. The plan should only update statistics that were automatically created and have not been updated for 30 days. The update should be based on all data in the table.You need to monitor query statistics and execution plans for SalesDb2.Which options should you set for the Query Store configuration? To answer, select the appropriate options in the answer area. Answer: Explanation: Operation Mode: Read-WriteIt is strongly recommended to activate size-based cleanup to makes sure that Query Store always runs in read-write mode and collects the latest data.Query Store Capture Mode: AutoAuto - Infrequent queries and queries with insignificant compile and execution duration are ignored. Thresholds for execution count, compile and runtime duration are internally determined.Size Based Cleanup Mode: AUTOIt is strongly recommended to activate size-based cleanup to makes sure that Query Store always runs in read-write mode and collects the latest data.OFF - size based cleanup won't be automatically activated.AUTO - size based cleanup will be automatically activated when size on disk reaches 90% of max_storage_size_mb. This is the default configuration value.Size based cleanup removes the least expensive and oldest queries first. It stops at approximately 80% of max_storage_size_mb.References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-database-query-store-options-transact-sql>

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store#set-the-optimal-query-capture-mode>QUESTION 185Drag and Drop QuestionNote: 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.You maintain a Microsoft SQL Server instance that contains the following databases SalesDb1, SalesDb2, and SalesDb3. Each database has table named Products and Sales. The following table shows the configuration of each database. The backup strategies for each database are described in the following table. Each full or differential backup operation writes into a new file and uses a different sequence number. You observe the following database corruption issues. SalesDb3 reports a number of database corruption issues related to error 823 and 824 when reading data pages. You must display the following information about the corrupted pages:- database name- impacted file id- impacted file physical name- impacted page id- event type that identifies the error type- error countUsers report performance issues when they run queries against SalesDb2. You plan to monitor query statistics and execution plans for SalesDb2 by using Query Store. The monitoring

strategy must meet the following requirements:- Perform automatic data cleanup when query store disk usage reaches 500 megabyte (MB).- Capture queries based on resource consumption.- Use a stale query threshold value of 60 days.The query optimizer generates suboptimal execution plans for a number of queries on the Sales table in SalesDb2. You will create a maintenance plan that updates statistics for the table. The plan should only update statistics that were automatically created and have not been updated for 30 days. The update should be based on all data in the table.Users reports that they encounter the following error when they query SalesDb1: "SQL Server detected a logical consistency-based I/O error: incorrect checksum (expected: 0x5d672d9b; actual: 0xdd672d98). It occurred during a read of page (1.232) in database ID 12 at offset 0x000000001d0000 in file F:Databases MSSQLServer DatabasesMSSQL13.MSSQL2016MSSQLDATASalesDb1.mdt."You must restore the impacted page from SalesDb1Full_1.bak. A single backup set named SalesDb1Log.bak was created since the latest full backup operation.You need to restore the impacted page. Which four Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order. Answer: Explanation: Step 1:Restore pageStart a page restore with a full database, file, or filegroup backup that contains the page. In the RESTORE DATABASE statement, use the PAGE clause to list the page IDs of all of the pages to be restored.Step 2:Restore log file with norecovery. Use the first file (FILE = 1).Step 3:Backup the tail-end of the log.Create a new log backup of the database that includes the final LSN of the restored pages, that is, the point at which the last restored page is taken offline.Step 4:Restore database with recovery. Use second file (FILE = 2).Restore the new log backup. After this new log backup is applied, the page restore is completed and the pages are now usable.Example:The following example restores four damaged pages of file B with NORECOVERY. Next, two log backups are applied with NORECOVERY, followed with the tail-log backup, which is restored with RECOVERY. This example performs an online restore. In the example, the file ID of file B is 1, and the page IDs of the damaged pages are 57, 202, 916, and 1016.RESTORE DATABASE <database> PAGE='1:57, 1:202, 1:916, 1:1016' FROM <file_backup_of_file_B>WITH NORECOVERY;RESTORE LOG <database> FROM <log_backup>WITH NORECOVERY;RESTORE LOG <database> FROM <log_backup>WITH NORECOVERY;BACKUP LOG <database> TO <new_log_backup>;RESTORE LOG <database> FROM <new_log_backup> WITH RECOVERY;References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/restore-pages-sql-server>QUESTION 186Hotspot 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.You maintain a Microsoft SQL Server instance that contains the following databases SalesDb1, SalesDb2, and SalesDb3. Each database has table named Products and Sales. The following table shows the configuration of each database. The backup strategies for each database are described in the following table. Each full or differential backup operation writes into a new file and uses a different sequence number. You observe the following database corruption issues. SalesDb3 reports a number of database corruption issues related to error 823 and 824 when reading data pages. You must display the following information about the corrupted pages:- database name- impacted file id- impacted file physical name- impacted page id- event type that identifies the error type- error countUsers report performance issues when they run queries against SalesDb2. You plan to monitor query statistics and execution plans for SalesDb2 by using Query Store. The monitoring strategy must meet the following requirements:- Perform automatic data cleanup when query store disk usage reaches 500 megabyte (MB).- Capture queries based on resource consumption.- Use a stale query threshold value of 60 days.The query optimizer generates suboptimal execution plans for a number of queries on the Sales table in SalesDb2. You will create a maintenance plan that updates statistics for the table. The plan should only update statistics that were automatically created and have not been updated for 30 days. The update should be based on all data in the table. Both SalesDb1 and SalesDb2 experience blocking and deadlock issues. You create an Extended Events session to monitor the databases. (Click the Exhibit button.)You need to create and configure the Extended Events session.Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic. Answer: Explanation: Lock escalation is the process of converting many fine-grain locks into fewer coarse-grain locks, reducing system overhead while increasing the probability of concurrency contention.References:

[https://technet.microsoft.com/en-us/library/ms184286\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms184286(v=sql.105).aspx)

<https://blobeater.blog/2017/02/06/using-extended-events-in-azure/>QUESTION 187Drag and Drop QuestionNote: 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. You maintain a Microsoft SQL Server instance that contains the following databases SalesDb1, SalesDb2, and SalesDb3. Each database has table named Products and Sales. The following table shows the configuration of each database. The backup strategies for each database are described in the following table. Each full or differential backup operation writes into a new file and uses a

different sequence number. You observe the following database corruption issues. SalesDb3 reports a number of database corruption issues related to error 823 and 824 when reading data pages. You must display the following information about the corrupted pages:- database name- impacted file id- impacted file physical name- impacted page id- event type that identifies the error type- error countUsers report performance issues when they run queries against SalesDb2. You plan to monitor query statistics and execution plans for SalesDb2 by using Query Store. The monitoring strategy must meet the following requirements:- Perform automatic data cleanup when query store disk usage reaches 500 megabyte (MB).- Capture queries based on resource consumption.- Use a stale query threshold value of 60 days.The query optimizer generates suboptimal execution plans for a number of queries on the Sales table in SalesDb2. You will create a maintenance plan that updates statistics for the table. The plan should only update statistics that were automatically created and have not been updated for 30 days. The update should be based on all data in the table. You need to write the query the maintenance plan will use to update the statistics.Which four Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order. Answer: Explanation:Step 1: Use SalesDB2?We specify the specific database SalesDB2, not the master.Step 2: Step 3:From scenario: The plan should only update statistics that were automatically created and have not been updated for 30 days.Step 4: Declare the cursor using WITH FULLSCAN.References:

<https://solutioncenter.apexsql.com/how-to-automate-and-schedule-sql-server-index-defragmentation/>QUESTION 188Hotspot QuestionNote: 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.You maintain a Microsoft SQL Server instance that contains the following databases SalesDb1, SalesDb2, and SalesDb3. Each database has table named Products and Sales. The following table shows the configuration of each database. The backup strategies for each database are described in the following table. Each full or differential backup operation writes into a new file and uses a different sequence number. You observe the following database corruption issues. SalesDb3 reports a number of database corruption issues related to error 823 and 824 when reading data pages. You must display the following information about the corrupted pages:- database name- impacted file id- impacted file physical name- impacted page id- event type that identifies the error type- error countUsers report performance issues when they run queries against SalesDb2. You plan to monitor query statistics and execution plans for SalesDb2 by using Query Store. The monitoring strategy must meet the following requirements:- Perform automatic data cleanup when query store disk usage reaches 500 megabyte (MB).- Capture queries based on resource consumption.- Use a stale query threshold value of 60 days.The query optimizer generates suboptimal execution plans for a number of queries on the Sales table in SalesDb2. You will create a maintenance plan that updates statistics for the table. The plan should only update statistics that were automatically created and have not been updated for 30 days. The update should be based on all data in the table.You need to recover SalesDb2 and SalesDb3 while minimizing the data loss.In the table below, identify the recovery strategy you should use for each database.NOTE: Make only one selection in each column. Each correct selection is worth one point. Answer: Explanation:SalesDB2:Row data are torn (error 824)How to Fix Torn PagesRun DBCC checkdb see for inconsistenciesCheck your error logs first and then restore your last backups and transaction logs.All backups are lost.SalesDB3: Checksum issues for data pages.FEPAIR_REBUILDPerforms 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. References: <https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/manage-the-suspect-pages-table-sql-server>

QUESTION 189Hotspot QuestionYou have a Microsoft SQL Server instance that hosts a database named DB1 that contains 800 gigabyte (GB) of data. The database is used 24 hours each day. You implement indexes and set the value of the Auto Update Statistics option set to True.Users report that queries take a long time to complete.You need to identify tables that meet the following requirements:- More than 1,000 rows have changed.- The statistics have not been updated in over a week.How should you complete the Transact-SQL statement? Answer: Explanation:Example:SELECT obj.name, obj.object_id, stat.name, stat.stats_id, last_updated, modification_counter FROM sys.objects AS objJOIN sys.stats stat ON stat.object_id = obj.object_idCROSS APPLY sys.dm_db_stats_properties(stat.object_id, stat.stats_id) AS sp WHERE modification_counter > 1000order by modification_counter desc;sys.sysindexes contains one row for each index and table in the current database.rowmodctr counts the total number of inserted, deleted, or updated rows since the last time statistics were updated for the table.Example 2:SELECTid AS [Table ID], OBJECT_NAME(id) AS [Table Name], name AS [Index Name], STATS_DATE(id, indid) AS [LastUpdated], rowmodctr AS [Rows Modified]FROM sys.sysindexesWHERE STATS_DATE(id, indid)<=DATEADD(DAY,-1,GETDATE())AND rowmodctr>10 AND (OBJECTPROPERTY(id,'IsUserTable'))=1References:

<https://social.msdn.microsoft.com/Forums/sqlserver/en-US/493b90e3-cdb8-4a16-8249-849ba0f82fcb/how-to-find-outdated-statistics-in-sql-server?forum=transactsql>QUESTION 190Drag and Drop QuestionYou administer a SQL Server instance. A database named

DB1 is corrupted. Backups of DB1 are available on a disk backup device located at Z:\Backups\Backup.bak. The backup device has the following backup sets:- a full database backup that is the first backup set on the device (FILE = 1)- a differential database backup that is the second backup set on the device (FILE = 2)- a transaction log backup that is the third backup set on the device (FILE = 3) You restore the full database backup and the differential database backup without rolling back the uncommitted transactions. You need to restore the transaction log backup and ensure the database is ready for use after restoring the transaction log. How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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: QUESTION 191 Hotspot Question You manage a database that uses the FULL recovery model and has a 60-minute recovery point objective (RPO). Full database backups are performed at 23:00 Monday through Friday. Each full database backup takes 6 hours to complete. Approximately 5% of the data changes each weekday. Backups must be performed as often as possible to reduce the number of files needed either a full recovery to any given time. You must minimize the duration of the backup tasks that are performed between 04:00 and 23:00. You need to reconfigure backups. Which time slot or schedule should you use for each backup type? Answer: QUESTION 192 Drag and Drop Question You manage a Microsoft SQL Server that has a database named salesOrders. Users connect to the database by using a client application. Users report that the application cannot connect to the database. You observe that the database storage has experienced a failure. You need to repair the database and ensure that applications can connect to the database. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. Answer: Explanation: The specified database must be in single-user mode to use one of the following repair options. 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: REPAIR_FAST Maintains syntax for backward compatibility only. No repair actions are performed. References:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql> QUESTION 193 Hotspot Question You view the Restore Transaction Log settings for a database named DB1 as shown in the following exhibit. Use the drop down menus to select the answer choice that completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point. Answer: Explanation: Log shipping secondary database can be kept in two restore modes: Restoring - This mode is also known as NORECOVERY mode. Standby - This mode is also known as Read-Only mode. When you use the NORECOVERY mode, the database will be in a restoring state and inaccessible to users, so uncommitted database transactions are not an issue. When you use the STANDBY mode, database will be in Read-Only state and users can access this database for read operations. Transactions in process on primary server or any uncommitted transaction cannot be read in secondary database. References:

<https://www.mssqltips.com/sqlservertip/3600/change-the-restore-mode-of-a-secondary-sql-server-database-in-log-shipping-with-ssms/>!!!RECOMMEND!!!1. | 2018 Latest 70-764 Exam Dumps (PDF & VCE) 332Q&As Download:

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

[YouTube.com/watch?v=50l_7j_PnRA](https://www.youtube.com/watch?v=50l_7j_PnRA)