

[2017 New 2017 Latest Updated 70-765 Dumps Free Download In Lead2pass (16-30)]

[2017 May Microsoft Official New Released 70-765 Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed!](#) The 70-765 braindumps are the latest, authenticated by expert and covering each and every aspect of 70-765 exam. Comparing with others, our exam questions are rich in variety. We offer PDF dumps and 70-765 VCE dumps. Welcome to choose. Following questions and answers are all new published by Microsoft Official Exam Center: <http://www.lead2pass.com/70-765.html>

QUESTION 16 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals. You manage a Microsoft SQL Server environment with several databases. You need to ensure that queries use statistical data and do not initialize values for local variables. Solution: You enable the QUERY_OPTIMIZER_HOTFIXES option for the databases. Does the solution meet the goal? A. Yes B. No Answer: B Explanation: QUERY_OPTIMIZER_HOTFIXES = { ON | OFF | PRIMARY } enables or disables query optimization hotfixes regardless of the compatibility level of the database. This is equivalent to Trace Flag 4199.

<https://msdn.microsoft.com/en-us/library/mt629158.aspx> QUESTION 17 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 manage on-premises and Microsoft Azure SQL Database instances for a company. Your environment must support the Microsoft SQL Server 2012 ODBC driver. You need to encrypt only specific columns in the database. What should you implement? A. transport-level encryption B. cell-level encryption C. Transparent Data Encryption D. Always Encrypted E. Encrypting File System F. BitLocker G. dynamic data masking Answer: D Explanation: To encrypt columns you can configure Always Encrypted. SQL Server Management Studio (SSMS) provides a wizard that helps you easily configure Always Encrypted by setting up the column master key, column encryption key, and encrypted columns for you. Always Encrypted allows client applications to encrypt sensitive data and never reveal the data or the encryption keys to SQL Server or Azure SQL Database. An Always Encrypted enabled driver, such as the ODBC Driver 13.1 for SQL Server, achieves this by transparently encrypting and decrypting sensitive data in the client application. Note: The ODBC driver automatically determines which query parameters correspond to sensitive database columns (protected using Always Encrypted), and encrypts the values of those parameters before passing the data to SQL Server or Azure SQL Database. Similarly, the driver transparently decrypts data retrieved from encrypted database columns in query results.

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault#encrypt-columns-configure-always-encrypted> [https://msdn.microsoft.com/en-us/library/mt637351\(v=sql.110\).aspx](https://msdn.microsoft.com/en-us/library/mt637351(v=sql.110).aspx) QUESTION 18 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. Your company has several Microsoft Azure SQL Database instances. Data encryption should be allowed to be implemented by the client applications that access the data. Encryption keys should not be made available to the database engine. You need to configure the database. What should you implement? A. transport-level encryption B. cell-level encryption C. Transparent Data Encryption D. Always Encrypted E. Encrypting File System F. BitLocker G. dynamic data masking Answer: A

Explanation: Using encryption during transit with Azure File Shares Azure File Storage supports HTTPS when using the REST API, but is more commonly used as an SMB file share attached to a VM. HTTPS is a transport-level security protocol. QUESTION 19 Drag and Drop Question You deploy a new Microsoft Azure SQL Database instance to support a variety of mobile applications and public websites. You plan to create a new security principal named User1. The principal must have access to select all current and future objects in a database named Reporting. The activity and authentication of the database user must be limited to the Reporting database. You need to create the new security principal. 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: Step 1, Step 2: First you need to create a login for SQL Azure, its syntax is as follows: CREATE LOGIN username WITH password='password'; This command needs to run in master db. Only afterwards can you run commands to create a user in the database. Step 3: Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a user: CREATE USER readonlyuser FROM LOGIN readonlylogin;

<https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/> QUESTION 20 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the

solution meets stated goals. You have a mission-critical application that stores data in a Microsoft SQL Server instance. The application runs several financial reports. The reports use a SQL Server-authenticated login named Reporting_User. All queries that write data to the database use Windows authentication. Users report that the queries used to provide data for the financial reports take a long time to complete. The queries consume the majority of CPU and memory resources on the database server. As a result, read-write queries for the application also take a long time to complete. You need to improve performance of the application while still allowing the report queries to finish. Solution: You configure the Resource Governor to limit the amount of memory, CPU, and IOPS used for the pool of all queries that the Reporting_user login can run concurrently. Does the solution meet the goal? A. Yes B. No

Answer: A Explanation: SQL Server Resource Governor is a feature that you can use to manage SQL Server workload and system resource consumption. Resource Governor enables you to specify limits on the amount of CPU, physical IO, and memory that incoming application requests can use. <https://msdn.microsoft.com/en-us/library/bb933866.aspx> QUESTION 21 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals. You have a mission-critical application that stores data in a Microsoft SQL Server instance. The application runs several financial reports. The reports use a SQL Server-authenticated login named Reporting_User. All queries that write data to the database use Windows authentication. Users report that the queries used to provide data for the financial reports take a long time to complete. The queries consume the majority of CPU and memory resources on the database server. As a result, read-write queries for the application also take a long time to complete. You need to improve performance of the application while still allowing the report queries to finish. Solution: You configure the Resource Governor to set the MAXDOP parameter to 0 for all queries against the database. Does the solution meet the goal? A. Yes B. No

Answer: B Explanation: SQL Server will consider parallel execution plans for queries, index data definition language (DDL) operations, and static and keyset-driven cursor population. You can override the max degree of parallelism value in queries by specifying the MAXDOP query hint in the query statement. [https://technet.microsoft.com/en-us/library/ms181007\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms181007(v=sql.105).aspx) QUESTION 22 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals. You have a mission-critical application that stores data in a Microsoft SQL Server instance. The application runs several financial reports. The reports use a SQL Server-authenticated login named Reporting_User. All queries that write data to the database use Windows authentication. Users report that the queries used to provide data for the financial reports take a long time to complete. The queries consume the majority of CPU and memory resources on the database server. As a result, read-write queries for the application also take a long time to complete. You need to improve performance of the application while still allowing the report queries to finish. Solution: You create a snapshot of the database. You configure all report queries to use the database snapshot. Does the solution meet the goal? A. Yes B. No

Answer: B Explanation: Use a Resource Governor instead. <https://msdn.microsoft.com/en-us/library/bb933866.aspx> QUESTION 23 A company has an on-premises Microsoft SQL Server 2014 environment. The company has a main office in Seattle, and remote offices in Amsterdam and Tokyo. You plan to deploy a Microsoft Azure SQL Database instance to support a new application. You expect to have 100 users from each office. In the past, users at remote sites reported issues when they used applications hosted at the Seattle office. You need to optimize performance for users running reports while minimizing costs. What should you do? A. Implement an elastic pool. B. Implement a standard database with readable secondaries in Asia and Europe, and then migrate the application. C. Implement replication from an on-premises SQL Server database to the Azure SQL Database instance. D. Deploy a database from the Premium service tier. Answer: B

Explanation: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-geo-replication-transact-sql#add-secondary-database> Case Study 1

- Automobile Parts Background You manage the Microsoft SQL Server environment for a company that manufactures and sells automobile parts. The environment includes the following servers: SRV1 and SRV2. SRV1 has 16 logical cores and hosts a SQL Server instance that supports a mission-critical application. The application has approximately 30,000 concurrent users and relies heavily on the use of temporary tables. The environment also includes the following databases: DB1, DB2, and Reporting. The Reporting database is protected with Transparent Data Encryption (TDE). You plan to migrate this database to a new server. You detach the database and copy it to the new server. You are performing tuning on a SQL Server database instance. The application which uses the database was written using an object relationship mapping (ORM) tool which maps tables as objects within the application code. There are 30 stored procedures that are regularly used by the application. QUESTION 24 Hotspot Question You need to resolve the identified issues. Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic. Answer: Explanation: From exhibit we see: Cost Threshold of Parallelism: 5 Optimize for Ad Hoc Workloads: false Max Degree of Parallelism: 0 (This is the default setting, which enables the server to determine the maximum degree of parallelism. It is fine.) Locks: 0 Query Wait: -1 Box 1: Optimize for Ad Hoc Workload Change the Optimize for Ad Hoc

Workload setting from false to 1/True. The optimize for ad hoc workloads option is used to improve the efficiency of the plan cache for workloads that contain many single use ad hoc batches. When this option is set to 1, the Database Engine stores a small compiled plan stub in the plan cache when a batch is compiled for the first time, instead of the full compiled plan. This helps to relieve memory pressure by not allowing the plan cache to become filled with compiled plans that are not reused.

QUESTION 25 Hotspot Question You need to optimize SRV1. What configuration changes should you implement? To answer, select the appropriate option from each list in the answer area. Answer: Explanation: From the scenario: SRV1 has 16 logical cores and hosts a SQL Server instance that supports a mission-critical application. The application has approximately 30,000 concurrent users and relies heavily on the use of temporary tables.

Box 1: Change the size of the tempdb log file. The size and physical placement of the tempdb database can affect the performance of a system. For example, if the size that is defined for tempdb is too small, part of the system-processing load may be taken up with autogrowing tempdb to the size required to support the workload every time you restart the instance of SQL Server. You can avoid this overhead by increasing the sizes of the tempdb data and log file.

Box 2: Add additional tempdb files. Create as many files as needed to maximize disk bandwidth. Using multiple files reduces tempdb storage contention and yields significantly better scalability. However, do not create too many files because this can reduce performance and increase management overhead. As a general guideline, create one data file for each CPU on the server (accounting for any affinity mask settings) and then adjust the number of files up or down as necessary.

Case Study 2 - Contoso, Ltd HOTSPOT Background You are the database administrator for Contoso, Ltd. The company has 200 offices around the world. The company has corporate executives that are located in offices in London, New York, Toronto, Sydney, and Tokyo. Contoso, Ltd. has a Microsoft Azure SQL Database environment. You plan to deploy a new Azure SQL Database to support a variety of mobile applications and public websites. The company is deploying a multi-tenant environment. The environment will host Azure SQL Database instances. The company plans to make the instances available to internal departments and partner companies. Contoso is in the final stages of setting up networking and communications for the environment. Existing Contoso and Customer instances need to be migrated to Azure virtual machines (VM) according to the following requirements:

- Contoso instances - should use the method requiring the least administrative effort to migrate instances to Azure Vms.
- Customer instances - should use a method that allows customers to bring their own licenses to Azure VMs. Customers have approved down time for the migration.

The company plans to deploy a new order entry application and a new business intelligence and analysis application. Each application will be supported by a new database. Contoso creates a new Azure SQL database named Reporting. The database will be used to support the company's financial reporting requirements. You associate the database with the Contoso Azure Active Directory domain. Each location database for the data entry application may have an unpredictable amount of activity. Data must be replicated to secondary databases in Azure datacenters in different regions. To support the application, you need to create a database named contosodb1 in the existing environment. Objects Database The contosodb1 database must support the following requirements:

- a size of at least 200 gigabytes (GB)
- 1,000 concurrent sessions
- point-in-time restore to any point in the two weeks prior to a failure
- minimize costs

Application For the business intelligence application, corporate executives must be able to view all data in near real-time with low network latency. Contoso has the following security, networking, and communications requirements:

- Multi-Location Load Balancing - to ensure customers have access to their tenants at multiple Azure locations across the world.
- Secure Message/Data Flow - to securely support communication between Azure and on-premises applications and services.
- Accounts should support accessing external domain resources and be configured in the most secure and lowest-maintenance way possible, including meeting the company policy of regular service account password changes.

QUESTION 26 Hotspot Question You need to configure the data entry and business intelligence databases. In the table below, identify the option that you must use for each database. NOTE: Make only one selection in each column. Answer: Explanation: Data Entry: Geo-replicated database only From Contoso scenario: Each location database for the data entry application may have an unpredictable amount of activity. Data must be replicated to secondary databases in Azure datacenters in different regions. Business intelligence: Elastic database pools only From Contoso scenario: For the business intelligence application, corporate executives must be able to view all data in near real-time with low network latency. SQL DB elastic pools provide a simple cost effective solution to manage the performance goals for multiple databases that have widely varying and unpredictable usage patterns. <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool>

QUESTION 27 Hotspot Question You need to create the contosodb1 database. How should you complete the Azure PowerShell command? To answer, select the appropriate Azure PowerShell segments in the answer area. Answer: Explanation: **Box 1:** New-AzureRmSqlDatabase New-AzureRmSqlDatabase creates a database or an elastic database. New-AzureRmSqlDatabase is a command with the Azure Resource Manager (AzureRM) module. Azure Resource Manager enables you to work with the resources in your solution as a group. **Case Study 3 - SQL Server Reporting** Background You manage a Microsoft SQL Server environment that includes the following databases: DB1, DB2, Reporting. The environment also includes SQL Reporting Services (SSRS) and SQL Server Analysis Services (SSAS). All SSRS

and SSAS servers use named instances. You configure a firewall rule for SSAS. Databases Database Name: DB1 Notes: This database was migrated from SQL Server 2012 to SQL Server 2016. Thousands of records are inserted into DB1 or updated each second. Inserts are made by many different external applications that your company's developers do not control. You observe that transaction log write latency is a bottleneck in performance. Because of the transient nature of all the data in this database, the business can tolerate some data loss in the event of a server shutdown. Database Name: DB2 Notes: This database was migrated from SQL Server 2012 to SQL Server 2016. Thousands of records are updated or inserted per second. You observe that the WRITELOG wait type is the highest aggregated wait type. Most writes must have no tolerance for data loss in the event of a server shutdown. The business has identified certain write queries where data loss is tolerable in the event of a server shutdown. Database Name: Reporting Notes: You create a SQL Server-authenticated login named BIAppUser on the SQL Server instance to support users of the Reporting database. The BIAppUser login is not a member of the sysadmin role. You plan to configure performance-monitoring alerts for this instance by using SQL Agent Alerts. QUESTION 28 Drag and Drop Question You create a login named BIAppUser. The login must be able to access the Reporting database. You need to grant access to the BIAppUser login in the database. How should you complete the Transact-SQL statements? 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: Explanation: Box 1: Reporting The user is to be created in the Reporting database. Box 2: CREATE USER Box 3: FOR LOGIN [BIAppUser] Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. Here is some sample Transact-SQL that creates a user: CREATE USER readonlyuser FROM LOGIN readonlylogin; <https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/> QUESTION 29 Hotspot Question You need to maximize performance of writes to each database without requiring changes to existing database tables. In the table below, identify the database setting that you must configure for each database. NOTE: Make only one selection in each column. Each correct selection is worth one point. Answer: Explanation: DB1: DELAYED_DURABILITY=FORCED From scenario: Thousands of records are inserted into DB1 or updated each second. Inserts are made by many different external applications that your company's developers do not control. You observe that transaction log write latency is a bottleneck in performance. Because of the transient nature of all the data in this database, the business can tolerate some data loss in the event of a server shutdown. With the DELAYED_DURABILITY=FORCED setting, every transaction that commits on the database is delayed durable. With the DELAYED_DURABILITY=ALLOWED setting, each transaction's durability is determined at the transaction level. Note: Delayed transaction durability reduces both latency and contention within the system because: * The transaction commit processing does not wait for log IO to finish and return control to the client. * Concurrent transactions are less likely to contend for log IO; instead, the log buffer can be flushed to disk in larger chunks, reducing contention, and increasing throughput. DB2: ALLOW_SNAPSHOT_ISOLATION ON and READ_COMMITTED_SNAPSHOT ON Snapshot isolation enhances concurrency for OLTP applications. Snapshot isolation must be enabled by setting the ALLOW_SNAPSHOT_ISOLATION ON database option before it is used in transactions. The following statements activate snapshot isolation and replace the default READ COMMITTED behavior with SNAPSHOT: ALTER DATABASE MyDatabase SET ALLOW_SNAPSHOT_ISOLATION ON ALTER DATABASE MyDatabase SET READ_COMMITTED_SNAPSHOT ON Setting the READ_COMMITTED_SNAPSHOT ON option allows access to versioned rows under the default READ COMMITTED isolation level. From scenario: The DB2 database was migrated from SQL Server 2012 to SQL Server 2016. Thousands of records are updated or inserted per second. You observe that the WRITELOG wait type is the highest aggregated wait type. Most writes must have no tolerance for data loss in the event of a server shutdown. The business has identified certain write queries where data loss is tolerable in the event of a server shutdown. <https://msdn.microsoft.com/en-us/library/dn449490.aspx> [https://msdn.microsoft.com/en-us/library/tcbchxcb\(v=vs.110\).aspx](https://msdn.microsoft.com/en-us/library/tcbchxcb(v=vs.110).aspx) QUESTION 30 Hotspot Question You need to set up the service accounts that the database engine and SQL Server Agent services will use. How should you design the solution? To answer, select the appropriate configuration options in the answer area. Answer: Explanation: Box 1: Domain Account The service startup account defines the Microsoft Windows account in which SQL Server Agent runs and its network permissions. SQL Server Agent runs as a specified user account. You select an account for the SQL Server Agent service by using SQL Server Configuration Manager, where you can choose from the following options: * Built-in account. You can choose from a list of the following built-in Windows service accounts: Local System account. * This account. Lets you specify the Windows domain account in which the SQL Server Agent service runs. Box 2: Domain users Microsoft recommends choosing a Windows user account that is not a member of the Windows Administrators group. Box 3: Managed Service Accounts When resources external to the SQL Server computer are needed, Microsoft recommends using a Managed Service Account (MSA), configured with the minimum privileges necessary. Note: A Managed Service Account (MSA) can run services on a computer in a secure and easy to maintain manner, while maintaining the capability to connect to network resources as a specific

user principal. <https://msdn.microsoft.com/en-us/library/ms191543.aspx> Microsoft 70-765 is often called the hardest of all Microsoft exams. Lead2pass helps you kill the Microsoft 70-765 exam challenge and achieve the perfect passing score with its latest practice test, packed into the revolutionary interactive VCE. This is the best way to prepare and pass the 70-765 exam. Microsoft 70-765 new questions on Google Drive: <https://drive.google.com/open?id=0B3Syig5i8gpDeGFILU5zdENPOFU> 2017 Microsoft 70-765 exam dumps (All 45 Q&As) from Lead2pass: <http://www.lead2pass.com/70-765.html> [100% Exam Pass Guaranteed]