

[Lead2pass New Free Share Lead2pass Microsoft 70-765 VCE Dumps With New Update Exam Questions (31-45)]

2017 October Microsoft Official New Released 70-765 Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed! How to 100% pass 70-765 exam? Lead2pass 70-765 dump is unparalleled in quality and is 100% guaranteed to make you pass 70-765 exam. All the 70-765 exam questions are the latest. Here are some free share of Microsoft 70-765 dumps. Following questions and answers are all new published by Microsoft Official Exam Center: <https://www.lead2pass.com/70-765.html>

QUESTION 31 Hotspot Question You need to open the firewall ports for use with SQL Server environment. In table below, identify the firewall port that you must use for each service. **NOTE:** Make only one selection in each column. **Answer:** **Explanation:** Report Server: 80 By default, the report server listens for HTTP requests on port 80. **QUESTION 32** 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 have deployed a GS-series virtual machine (VM) in Microsoft Azure. You plan to deploy Microsoft SQL Server. You need to deploy a 30 megabyte (MB) database that requires 100 IOPS to be guaranteed while minimizing costs. Which storage option should you use? **A.** Premium P10 disk storage **B.** Premium P20 disk storage **C.** Premium P30 disk storage **D.** Standard locally redundant disk storage **E.** Standard geo-redundant disk storage **F.** Standard zone redundant blob storage **G.** Standard locally redundant blob storage **H.** Standard geo-redundant blob storage **Answer:** A **Explanation:** Premium Storage Disks Limits When you provision a disk against a Premium Storage account, how much input/output operations per second (IOPS) and throughput (bandwidth) it can get depends on the size of the disk. Currently, there are three types of Premium Storage disks: P10, P20, and P30. Each one has specific limits for IOPS and throughput as specified in the following table:

<https://docs.microsoft.com/en-us/azure/storage/storage-premium-storage> **QUESTION 33** You plan to migrate a database To Microsoft Azure SQL Database. The database requires 500 gigabytes (GB) of storage. The database must support 50 concurrent logins. You must minimize the cost associated with hosting the database. You need to create the database. Which pricing tier should you use? **A.** Standard S3 pricing tier **B.** Premium P2 tier **C.** Standard S2 pricing tier **D.** Premium P1 tier **Answer:** D **Explanation:** For a database size of 500 GB the Premium tier is required. Both P1 and P2 are adequate. P1 is preferred as it is cheaper. **Note:** **QUESTION 34** 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. Your company plans to use Microsoft Azure Resource Manager templates for all future deployments of SQL Server on Azure virtual machines. You need to create the templates. **Solution:** You create the desired SQL Server configuration in an Azure Resource Group, then export the Resource Group template and save it to the Templates Library. Does the solution meet the goal? **A.** Yes **B.** No **Answer:** B **Explanation:** Azure Resource Manager template consists of JSON, and expressions that you can use to construct values for your deployment. A good JSON editor, not a Resource Group template, can simplify the task of creating templates. **Note:** In its simplest structure, a Azure Resource Manager template contains the following elements: {"\$schema": "<http://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#>", "contentVersion": "", "parameters": { }, "variables": { }, "resources": [], "outputs": { } }

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-authoring-templates> **QUESTION 35** 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 have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in a development environment. Each VM has a dedicated disk for backups. You need to backup a database to the local disk on a VM. The backup must be replicated to another region. Which storage option should you use? **A.** Premium P10 disk storage **B.** Premium P20 disk storage **C.** Premium P30 disk storage **D.** Standard locally redundant disk storage **E.** Standard geo-redundant disk storage **F.** Standard zone redundant blob storage **G.** Standard locally redundant blob storage **H.** Standard geo-redundant blob storage **Answer:** E **Explanation:** Note: SQL Database automatically creates a database backups and uses Azure read- access geo-redundant storage (RA-GRS) to provide geo-redundancy. These backups are created automatically and at no additional charge. You don't need to do anything to make them happen. Database backups are an essential part of any business continuity and disaster recovery strategy because they protect your data from accidental corruption or deletion.

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-automated-backups> **QUESTION 36** 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. Your company plans to use Microsoft Azure Resource Manager templates for all future deployments of SQL Server on Azure virtual machines. You need to create the templates. **Solution:** You use Visual Studio to create a JSON template

that defines the deployment and configuration settings for the SQL Server environment. Does the solution meet the goal? A. Yes B. No

Answer: A
Explanation: Azure Resource Manager template consists of JSON, not XAML, and expressions that you can use to construct values for your deployment. A good JSON editor can simplify the task of creating templates. Note: In its simplest structure, an Azure Resource Manager template contains the following elements: {"\$schema": "

<http://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#>","contentVersion": "", "parameters": { }, "variables": { }, "resources": [], "outputs": { } }

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-authoring-templates> QUESTION 37 You plan to deploy an on-premises SQL Server 2014 database to Azure SQL Database. You have the following requirements: - Maximum database size of 500 GB - A point-in-time-restore of 35 days - Maximum database transaction units (DTUs) of 500 You need to choose the correct service tier and performance level. Which service tier should you choose? A. Standard S3 B. Premium P4 C. Standard S0 D. Basic
Answer: B
Explanation: You should choose Premium P4. The Premium tier is the highest Azure SQL Database tier offered. This tier is used for databases and application that require the highest level of performance and recovery. The P4 level supports a maximum of 500 DTUs, a maximum database size of 500 GB, and a point-in-time-restore to any point in the last 35 days. You should not choose Basic. The Basic service tier only supports a maximum database size of 2 GB, a maximum of 5 DTUs, and a point-in-time-restore of 7 days. You should not choose Standard S0. The Standard tier with a performance level of S0 only supports a maximum database size of 250 GB, a maximum of 10 DTUs, and a point-in-time-restore of 35 days. You should not choose Standard S3. The Standard tier with a performance level of S3 only supports a maximum database size of 250 GB, a maximum of 100 DTUs, and a point-in-time-restore of 35 days.

QUESTION 38 You are the database administrator for your company. Your company has one main office and two branch offices. You plan to create three databases named DB1, DB2, and DB3 that will be hosted on one Azure SQL Database server. You have the following requirements: - The main office must be able to connect to all three databases. - The branch offices must be able to connect to DB2 and DB3. - The branch offices must not be able to access DB1. You need to configure transparent data encryption (TDE) for DB1. Which two actions should you perform? Each correct answer presents part of the solution. A. Run CREATE CERTIFICATE cert1 WITH Subject = TDE Cert1 on DB1. B. Connect to DB1. C. Run ALTER DATABASE DB1 SET ENCRYPTION ON;. D. Connect to the master database. E. Run CREATE MASTER KEY on the master database.
Answer: B C
Explanation: You should connect to DB1. To encrypt DB1, you connect directly to DB1. When you connect to DB1, you use your dbmanager or administrative credentials. You should run ALTER DATABASE DB1 SET ENCRYPTION ON. You use the ALTER DATABASE DB1 SET ENCRYPTION ON statement to encrypt the database. This is the statement that turns on TDE for Azure SQL Database. You should not run CREATE MASTER KEY on the master database. You would execute CREATE MASTER KEY when you have to encrypt an on-premises database. You should not run CREATE CERTIFICATE cert1 WITH Subject = TDE Cert1 on DB1. You do not need to create a certificate to encrypt an Azure SQL database. This would be a part of the solution when you encrypt an on-premises database. You should not connect to the master database. To encrypt DB1, you need to connect directly to DB1, not to the master database.

QUESTION 39 You manage an on-premises, multi-tier application that has the following configuration: - Two SQL Server 2012 databases named SQL1 and SQL2 - Two application servers named AppServer1 and AppServer2 that run IIS You plan to move your application to Azure. You need to ensure that during an Azure update cycle or a hardware failure, the application remains available. Which two deployment configurations should you implement? Each correct answer presents part of the solution. A. Deploy AppServer1 and AppServer2 in a single availability set. B. Deploy all servers in a single availability set. C. Deploy SQL1 and AppServer1 in a single availability set. D. Deploy SQL2 and AppServer2 in a single availability set. E. Deploy SQL1 and SQL2 in a single availability set.
Answer: A E
Explanation: You should deploy SQL1 and SQL2 in a single availability set. Using availability sets allows you to build in redundancy for your Azure services. By grouping related virtual machines and services (tiers) into an availability set (in this case, deploying both of your databases into an availability set), you ensure that if there is a planned or unplanned outage, your services will remain available. At the most basic level, virtual machines in an availability set are put into a different fault domain and update domain. An update domain allows virtual machines to have updates installed and then the virtual machines are rebooted together. If you have two virtual machines in an availability set, each in its own update domain, a rebooting of one server does not bring down all of the servers in a given tier. A fault domain operates in the same manner, so if there is a physical problem with a server, rack, network, or other service, both machines are separated, and services will continue. You should deploy AppServer1 and AppServer2 in a single availability set. Using availability sets allows to build in redundancy for your Azure services. By grouping related virtual machines and services (tiers) an availability set (in this case, deploying both of your application servers into an availability set), ensure that if there is a planned or unplanned outage, your services will remain available. You should not deploy SQL1 and AppServer1 in a single availability set. Microsoft recommends grouping similar sets of services in an availability set. This will help

ensure that each tier of your application is protected by an availability set. SQL1 and AppServer1 provide different services, so they should not be grouped together. You should not deploy SQL2 and AppServer2 in a single availability set. SQL2 and AppServer2 provide different services, so they should not be grouped together. You should not deploy all servers in a single availability set. This will not provide the fault tolerance needed, as all machines would be rebooted (or suffer a hardware failure) together.

QUESTION 40 You are the database administrator in your company. You plan to create 10 identical environments that use SQL Server 2016 as a database engine. Each environment has the following custom requirements: - Three user databases must be preinstalled. - The tempdb database must contain eight data files that are 1024 MB each. - Trace flag 2371 must be turned at the instance level. The solution must meet the following requirements: - The instance must be preconfigured. - No other database features are required in the future. - The solution must use the minimum administrative effort. You need to prepare the environments. What should you do?

A. Provision 10 Azure virtual machines that each contain SQL Server 2016, installed by using the default settings. B. Create an installation configuration file and perform unattended installations of SQL Server 2016. C. Create a virtual machine template by using a prepared instance of SQL Server 2016. D. Create a virtual machine template by using a complete instance of SQL Server 2016.

Answer: D

Explanation: You should create a virtual machine template by using a complete instance of SQL Server 2016. You use the sysprep tool to prepare a complete instance of SQL Server 2016. By using a complete instance, SQL Server, the network, and the users are all created, and the system cannot be reconfigured during the installation process. You should not create a virtual machine template by using a prepared instance of SQL Server 2016. A virtual machine template that contains a prepared instance of SQL Server could be modified during the installation process. For example, some features could be added or removed during the installation process. The prepared instance contains a preconfigured version of SQL Server without network and user configurations. You should not create an installation configuration file and perform unattended installations of SQL Server 2016. This could be a valid option, but you would need to prepare 10 installations, and this would require a lot of effort. There are also some other administrative actions that must be performed that could also increase the amount of effort required, such as user database configuration. You should not provision 10 Azure virtual machines that each contain SQL Server 2016, installed by using the default settings. A virtual machine created by using the Azure library provides you with a standard installation, and the requirements call for a custom installation. The instance will need some specific parameters for your custom installation. Therefore, you would have to customize each Azure virtual machine.

QUESTION 41 You have a server named Server1 that is hosted in an Azure virtual machine. Server1 contains the following: - One instance of SQL Server 2016 Enterprise - 10 databases - 500 stored procedures You have a database named Database1 that is hosted on Server1. Database1 contains 100 queries that are executed dynamically from web applications. You plan to remove data from the procedure cache on Database1. You have the following requirements: - Changes to Database1 must not affect other databases that are hosted on Server1 - Changes to Database1 must not affect the performance of queries that are stored in other databases. - The solution must minimize administrative effort. You need to remove the data from the procedure cache as quickly as possible. What should you do?

A. Run DBCC FREEPROCCACHE. B. Run ALTER DATABASE SCOPED CONFIGURATION CLEAR PROCEDURE CACHE in the context of Database 1. C. Run DBCC DROPCLEANBUFFERS. D. Write a script that iterates through each stored procedure definition and add WITH RECOMPILE to the definition.

Answer: B

Explanation: You should run ALTER DATABASE SCOPED CONFIGURATION CLEAR PROCEDURE CACHE in the context of Database1. This statement lets you change the settings of a database without affecting other databases that are installed on the instance of SQL Server 2016. You should not run DBCC FREEPROCCACHE. DBCC FREEPROCCACHE would clean the entire plan cache and would affect all databases. It is possible to remove a single plan from the cache by using the plan_handle argument to DBCC FREEPROCCACHE, but you would have to identify all plans that are related to Database1, which requires a lot more administrative effort. You should not run DBCC DROPCLEANBUFFERS. DBCC DROPCLEANBUFFERS will remove the clean pages from the buffer cache. Columnstore pages are removed from the columnstore cache. You should not write a script that will iterate through each stored procedure definition and add WITH RECOMPILE to the definition. Each time the procedure is called, it will be recompiled, and this might degrade the server's performance. This approach would require additional administrative effort to produce the script. In addition, some calls are made from the web application and the script would not have any control over these calls. In earlier versions of SQL Server, prior to the availability of the ALTER DATABASE statement, this option would have been the way to avoid affecting other databases.

QUESTION 42 You are designing a Windows Azure SQL Database for an order fulfillment system. You create a table named Sales.Orders with the following script. Each order is tracked by using one of the following statuses: - Fulfilled - Shipped - Ordered - Received You need to design the database to ensure that that you can retrieve the following information: - The current status of an order - The previous status of an order. - The date when the status changed. - The solution must minimize storage. More than one answer choice may achieve the goal. Select the BEST answer.

A. To the Sales.Orders table, add three columns named Status, PreviousStatus and ChangeDate. Update rows as the order status

changes.B. Create a new table named Sales.OrderStatus that contains three columns named OrderID, StatusDate, and Status. Insert new rows into the table as the order status changes.C. Implement change data capture on the Sales.Orders table.D. To the Sales.Orders table, add three columns named FulfilledDate, ShippedDate, and ReceivedDate. Update the value of each column from null to the appropriate date as the order status changes. Answer: B QUESTION 43You are using dynamic management views to monitor an SQL Server server named SQL1.A database administrator named Dbal1 must monitor the health of SQL1.You need to ensure that Dbal1 can access dynamic management views for SQL1. The solution must use the principle of least privilege.Which permissions should you assign to Dbal1? A. VIEW ANY DEFINITIONB. VIEW SERVER STATEC. VIEW DEFINITIOND. CONTROL SERVER Answer: BExplanation: To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission. QUESTION 44You are the administrator for a SQL Server 2016 instance that stores the data for an online transaction processing sales system.The company takes full backups every week; differential backups on the days with no full backups; and hourly transaction backups. These backups are stored on a backup server in the company's data center.Every week, the company places the full backup on a tape and sends it to a third-party backup storage system.The company is worried that a disaster might occur that could destroy their computer center and cause them to lose orders.You need to determine the best method for providing the smallest amount of data loss and downtime without leasing or purchasing additional physical locations.What should you do? More than one answer choice may achieve the goal. Select the BEST answer. A. Set up SQL Server Always On with a SQL Azure database as a replica.B. Set up SQL Server Always On by using a SQL Server on a Windows Azure Virtual Machine.C. Put the differential backup on tape and send it to the third-party backup storage system.D. Use the Microsoft SQL Server Backup to Microsoft Windows Azure Tool to direct all backups to a different geographical location. Answer: D QUESTION 45You have a SQL Server 2016 database named DB1.You plan to import a large number of records from a SQL Azure database to DB1.You need to recommend a solution to minimize the amount of space used in the transaction log during the import operation.What should you include in the recommendation? A. The bulk-logged recovery modelB. The full recovery modelC. A new partitioned tableD. A new log fileE. A new file group Answer: AExplanation: Compared to the full recovery model, which fully logs all transactions, the bulk-logged recovery model minimally logs bulk operations, although fully logging other transactions. The bulk-logged recovery model protects against media failure and, for bulk operations, provides the best performance and least log space usage.Note:The bulk-logged recovery model is a special-purpose recovery model that should be used only intermittently to improve the performance of certain large-scale bulk operations, such as bulk imports of large amounts of data. Always up-to-date Lead2pass 70-765 VCE - everything you need for your Microsoft 70-765 exam to pass. Our Microsoft 70-765 software allows you to practise exam dumps in real 70-765 exam environment. Welcome to choose. More Microsoft 70-765 new questions (with images) on Google Drive:
<https://drive.google.com/open?id=0B3Syig5i8gpDejczWp0aURaSnM> 2017 Microsoft 70-765 exam dumps (All 115 Q&As) from Lead2pass: <https://www.lead2pass.com/70-765.html> [100% Exam Pass Guaranteed]