DUMPSBOSS.

Provisioning SQL Databases

Microsoft 70-765

Version Demo

Total Demo Questions: 15

Total Premium Questions: 223

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Implementing SQL in Azure	72
Topic 2, Manage databases and instances	72
Topic 3, Manage Storage	79
Total	223

QUESTION NO: 1

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 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, 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": { }

}

References: https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-authoring-templates

QUESTION NO: 2

You have a database named DB1 that is 3 TB. DB1 contains a fact table that is 1.2 TB.

You load 200 GB of new data to the fact table from a line-of-business application.

Users of DB1 notice that reports render more slowly after you loaded the data.

What are two possible causes of the performance issue? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. transaction log growth

B. out-of-date statistics

C. page corruption

D. index fragmentation

ANSWER: B D

Explanation:

Reference: https://www.sqlshack.com/gathering-sql-server-indexes-statistics-and-usage-information/

QUESTION NO: 3

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 AppServerl 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 AppServerl and AppServer2 in a single availability set.

B. Deploy all servers in a single availability set.

C. Deploy SQL1 and AppServerl 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 AppServerl and AppServer2 in a single availability set.

You should deploy SQL1 and SQL2 in a single availability set.

Note: 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.

Incorrect Answers:

B: 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.

C: SQL1 and AppServerl provide different services, so they should not be grouped together.

D: 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.

References: https://docs.microsoft.com/en-us/azure/virtual-machines/windows/manage-availability

QUESTION NO: 4 - (DRAG DROP)

DRAG DROP

You have a fact table named FactSales that is 100 GB. FactSales is in a data warehouse that is partitioned by month.

You discover that queries perform index scan operations when the queries should perform index seek operations.

You need to optimize the query plan to reduce the number of scans.

How should you complete the statement? To answer, drag the appropriate values to the correct targets. Each value 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.

Select and Place:

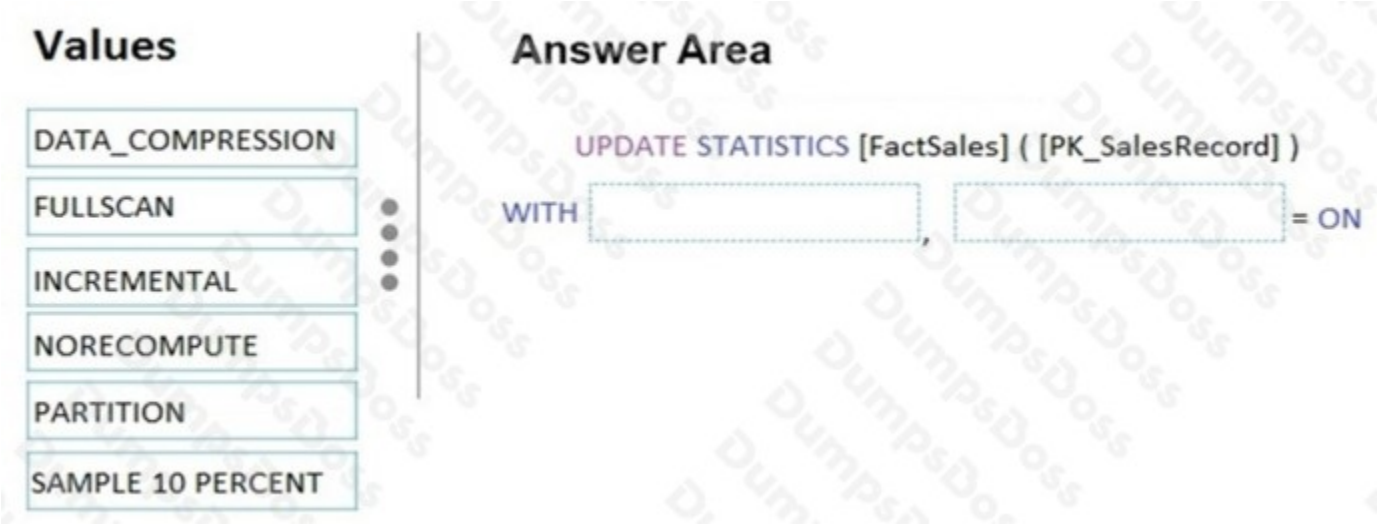

ANSWER:

Explanation:

References: https://docs.microsoft.com/en-us/sql/t-sql/statements/update-statistics-transact-sql?view=sql-server-2017

QUESTION NO: 5

You create an availability group that has replicas named HA/Server01 and HA/Server02.

Currently, HA/Server01 is the primary replica.

You have multiple queries that read data and produce reports from the database.

You need to offload the reporting workload to the secondary replica when HA/Server01 is the primary replica.

What should you do?

A. Set the Availability Mode property of HA/Server02 to Asynchronous commit.

B. Set the Readable Secondary property of HA/Server02 to Read-intent only.

C. Set the Connections in Primary Role property of HA/Server01 to Allow read/write connections.

D. Set the Availability Mode property of HA/Server01 to Asynchronous commit.

ANSWER: B

Explanation:

To set up a readable secondary replica, you first create an availability group. Then you add replicas. You can choose either Yes or Read-intent only options.

References: http://msdn.microsoft.com/en-us/library/jj542414.aspx

QUESTION NO: 6

You plan to deploy Microsoft SQL Server on a Microsoft Azure virtual machine. The virtual machine will have two databases. Each database will reside on a separate VHD and will be between 600 and 800 GB.

Each database will have the I/O requirements shown in the following table.

Database name	Maximum IOPS
DB1	4,000
DB2	1,200

You are evaluating whether to use the P30 storage disk type.

What is the minimum number of disks required for each database when using P30 storage disk type? (Select two.)

A. DB1: 0

B. DB1: 1

C. DB1: 2

D. DB1: 3

E. DB1: 4

F. DB2: 0

G. DB2: 1

H. DB2: 2

I. DB2: 3

J. DB2: 4

ANSWER: C G

Explanation:

P30 stats: Disk size is 1024 GB (1 TB), IOPS per disk is 5000.

Recommendation: Use a minimum of 2 P30 disks (1 for log files and 1 for data files and TempDB; or stripe two or more disks and store all files in a single volume).

References: https://docs.microsoft.com/en-us/azure/virtual-machines/windows/premium-storage#scalability-and-performance-targets

QUESTION NO: 7

You have an on-premises Microsoft SQL Server instance.

You create a database named DB1.

You need to configure DB1 to support In-Memory OLTP.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Enable In-Memory OLTP for the instance.

B. Configure a new file group for memory-optimized data.

C. Enable FILESTREAM for the instance.

D. Configure a new file for a file group.

E. Configure a new file group for FILESTREAM data.

ANSWER: A B

Explanation:

References:

https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/survey-of-initial-areas-in-in-memory-oltp?view=sql-server-ver15

QUESTION NO: 8 - (HOTSPOT)

HOTSPOT

You plan to deploy two new Microsoft Azure SQL Database instances. Once instance will support a data entry application. The other instance will support the company's business intelligence efforts. The databases will be accessed by mobile applications from public IP addresses.

You need to ensure that the database instances meet the following requirements:

- The database administration team must receive alerts for any suspicious activity in the data entry database, including potential SQL injection attacks.
- Executives around the world must have access to the business intelligence application.
- Sensitive data must never be transmitted. Sensitive data must not be stored in plain text in the database.

In the table below, identify the feature that you must implement for each database.

NOTE: Make only one selection in each column. Each correct selection is work one point.

Hot Area:

ANSWER:

Explanation:

Data entry: Threat Detection

SQL Threat Detection provides a new layer of security, which enables customers to detect and respond to potential threats as they occur by providing security alerts on anomalous activities. Users receive an alert upon suspicious database activities, potential vulnerabilities, and SQL injection attacks, as well as anomalous database access patterns.

Business intelligence: Dynamic Data Masking

Dynamic data masking limits (DDM) sensitive data exposure by masking it to non-privileged users. It can be used to greatly simplify the design and coding of security in your application.

References:

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-threat-detection https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking

QUESTION NO: 9

You administer a Microsoft SQL Server 2012 server. You plan to deploy new features to an application.

You need to evaluate existing and potential clustered and non-clustered indexes that will improve performance.

What should you do?

A. Query the sys.dm_db_index_usage_stats DMV.

B. Query the sys.dm_db_missing_index_details DMV.

C. Use the Database Engine Tuning Advisor.

D. Query the sys.dm_db_missing_index_columns DMV.

ANSWER: C

Explanation:

The Microsoft Database Engine Tuning Advisor (DTA) analyzes databases and makes recommendations that you can use to optimize query performance. You can use the Database Engine Tuning Advisor to select and create an optimal set of indexes, indexed views, or table partitions without having an expert understanding of the database structure or the internals of SQL Server.

Incorrect Answers:

A: sys.dm_db_index_usage_stats returns counts of different types of index operations and the time each type of operation was last performed in SQL Server.

B: sys.dm_db_missing_index_details returns detailed information about missing indexes, excluding spatial indexes.

D: sys.dm_db_missing_index_columns returns information about database table columns that are missing an index, excluding spatial indexes.

References: https://docs.microsoft.com/en-us/sql/relational-databases/performance/database-engine-tuning-advisor

QUESTION NO: 10 - (HOTSPOT)

HOTSPOT

You have an on-premises database.

You plan to migrate the database to Microsoft SQL Server on a Microsoft Azure virtual machine.

You move the database files to Azure.

You need to attach the database files to the SQL Server instance on the virtual machine. The solution must ensure that you can run file snapshot backups.

How should you complete the statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

ANSWER:

Explanation:

References: https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-sql-server-transact-sql

QUESTION NO: 11

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.

You are tuning the performance of a virtual machine that hosts a Microsoft SQL Server instance.

The virtual machine originally had four CPU cores and now has 32 CPU cores.

The SQL Server instance uses the default settings and has an OLTP database named db1. The largest table in db1 is a key value store table named table1.

Several reports use the PIVOT statement and access more than 100 million rows in table1.

You discover that when the reports run, there are PAGELATCH_IO waits on PFS pages 2:1:1, 2:2:1, 2:3:1, and 2:4:1 within the tempdb database.

You need to prevent the PAGELATCH_IO waits from occurring.

Solution: You add more files to db1.

Does this meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

From SQL Server's perspective, you can measure the I/O latency from sys.dm_os_wait_stats. If you consistently see high waiting for PAGELATCH_IO, you can benefit from a faster I/O subsystem for SQL Server.

A cause can be poor design of your database - you may wish to split out data located on 'hot pages', which are accessed frequently and which you might identify as the causes of your latch contention. For example, if you have a currency table with a data page containing 100 rows, of which 1 is updated per transaction and you have a transaction rate of 200/sec, you could see page latch queues of 100 or more. If each page latch wait costs just 5ms before clearing, this represents a full half-second delay for each update. In this case, splitting out the currency rows into different tables might prove more performant (if less normalized and logically structured).

References: https://www.mssqltips.com/sqlservertip/3088/explanation-of-sql-server-io-and-latches/

QUESTION NO: 12

You have Microsoft SQL Server on a DS-series Microsoft Azure virtual machine. The virtual machine has 28 GB of memory.

You discover the following performance statistics on the server:

- The average Page life expectancy is 30.
- The server has excessive PAGELATCH_IO waits.

You need to decrease the PAGELATCH_IO waits.

What should you do?

A. Enable large-page support.

B. Enable Lock pages in memory.

C. Configure buffer pool extensions.

D. Add more tempdb files.

ANSWER: D

Explanation:

Reference: https://www.brentozar.com/archive/2014/05/tell-need-tempdb-files/

QUESTION NO: 13

You have the databases configured as shown in the following table.

Database name	Location	Platform
DB1	Microsoft Azure	Microsoft Azure SQL Database
DB2	Microsoft Azure	Microsoft SQL Server 2016 on a Microsoft Azure virtual machine
DB3	On-premises	Microsoft SQL Server 2016
DB4	On-premises	Microsoft SQL Server 2014

Which two databases can use the Stretch Database feature? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. DB1

B. DB2

C. DB3

D. DB4

ANSWER: A C

Explanation:

References: https://docs.microsoft.com/en-us/sql/sql-server/stretch-database/enable-stretch-database-for-a-database?view=sql-server-2017

QUESTION NO: 14 - (DRAG DROP)

DRAG DROP

You have a Microsoft SQL Server instance which hosts all of your corporate databases. A database named Sales stores information about customers and their contact information.

You use the following processes for backing up the database:

- All databases are configured for full recovery model.
- Full backups are performed every morning at 2:00 AM.
- Log backups are performed every hour starting at 9:00 AM.

At 9:35 AM, a member of the sales team mistakenly updates all customer records.

You need to recover the database to a stable state and recover as much data as possible without recovering the changes that the sales team member made.

Which four 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.

Select and Place:

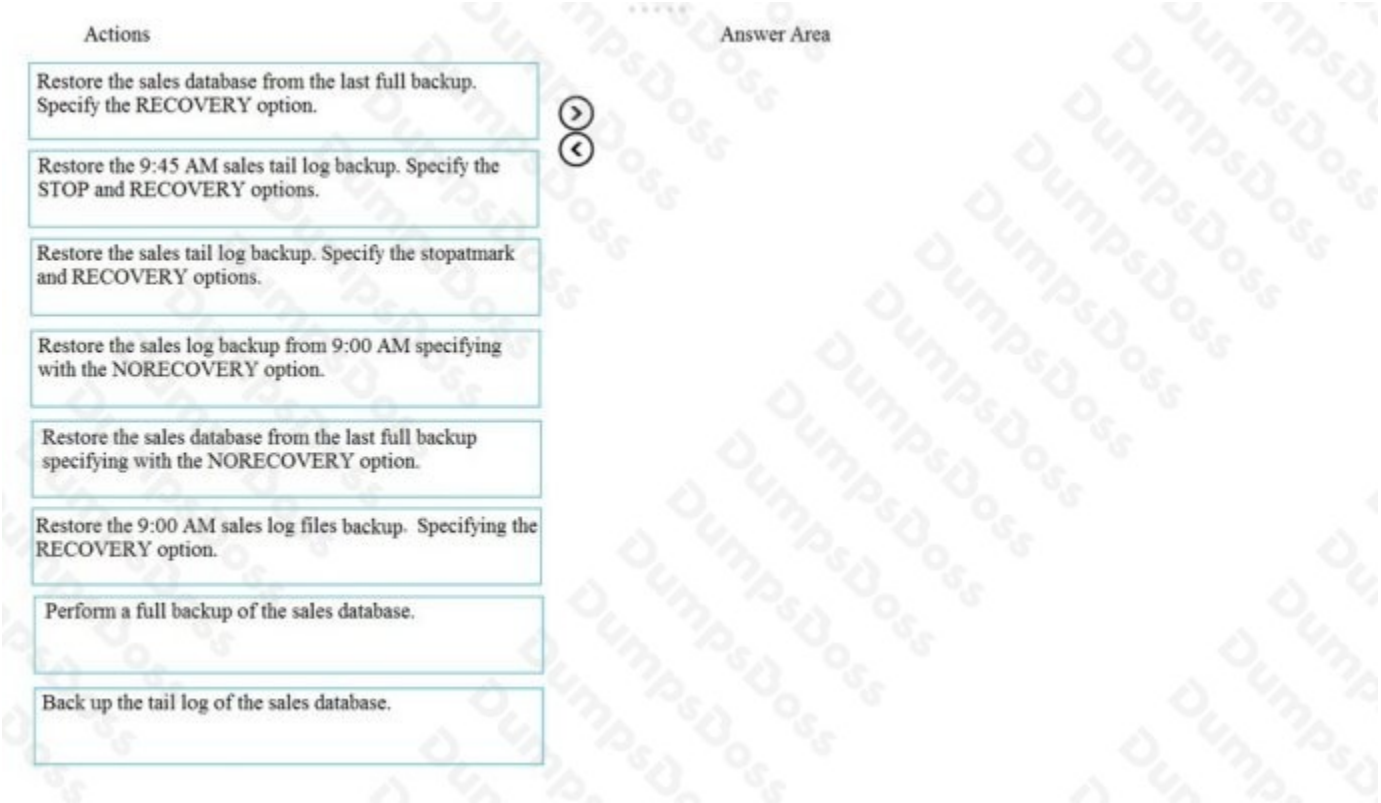

ANSWER:

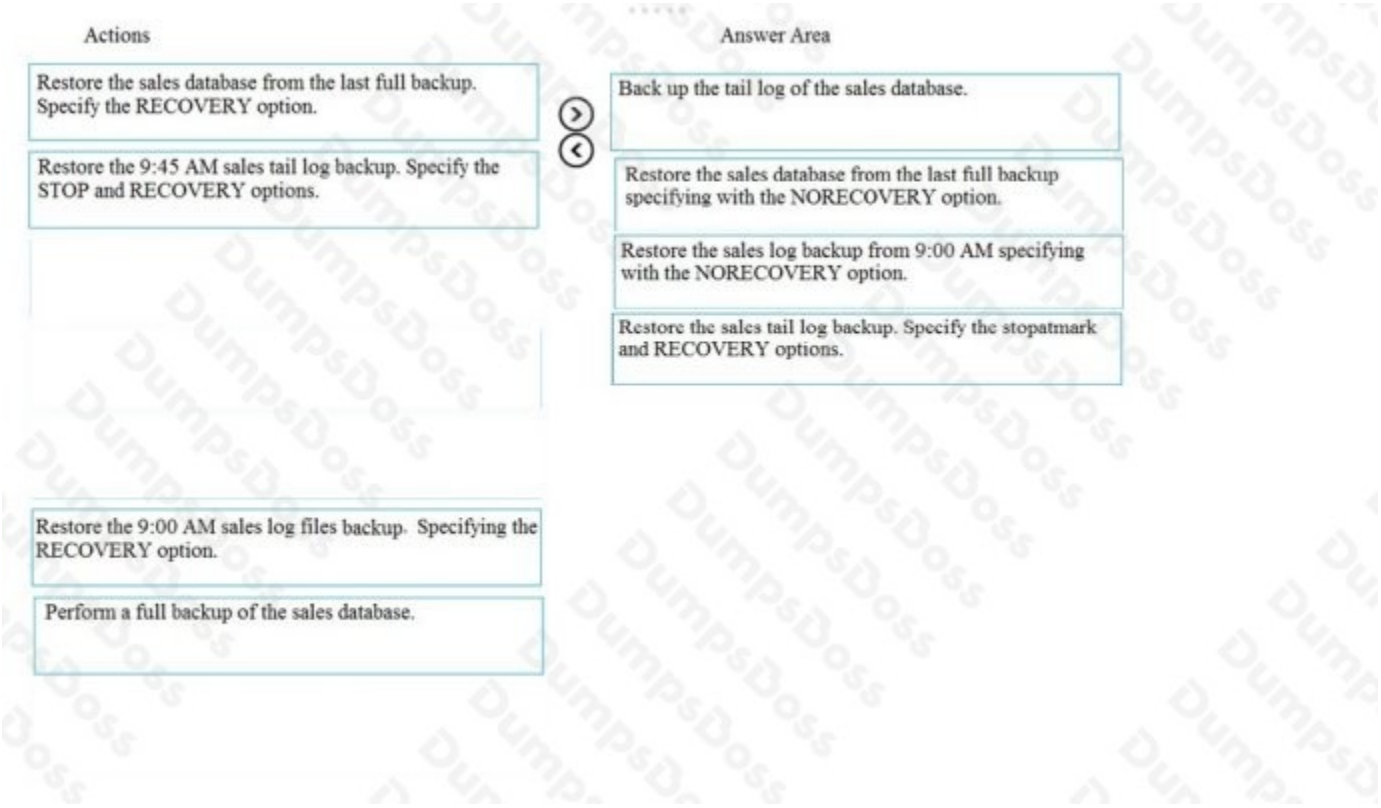

Explanation:

QUESTION NO: 15

You have Microsoft SQL Server on a Microsoft Azure virtual machine.

You have two Windows accounts named ServiceAccount1 and ServiceAccount2. The SQL Server Agent runs as ServiceAccount1.

You need to run SQL Server Agent job steps by using ServiceAccount2.

Which cmdlet should you run first?

A. Set-ADServiceAccount

B. Set-SqlCredential

C. New-ADServiceAccount

D. New-SqlCredential

ANSWER: C

Explanation:

The New-ADServiceAccount command creates a new Active Directory managed service account or group managed service account object.

Incorrect Answers:

A: The Set-ADServiceAccount cmdlet modifies the properties of an Active Directory managed service account (MSA). You can modify commonly used property values by using the cmdlet parameters.

B: The Set-SqlCredential cmdlet sets the Identity and password properties for a SQL Credential object using this cmdlet.

D: The New-SqlCredential cmdlet creates a new SQL Server credential object. A SQL Server credential object is used to store authentication information. The SQL Server credential is required when backing up to or restoring from the Windows Azure storage service, and is used to store the Windows Azure storage account name and access key information.

References: https://docs.microsoft.com/en-us/powershell/module/addsadministration/new-adserviceaccount?view=win10-ps