



DP-300^{Q&As}

Administering Relational Databases on Microsoft Azure

Pass Microsoft DP-300 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/dp-300.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

DRAG DROP

You are creating a managed data warehouse solution on Microsoft Azure.

You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails.

You need to configure Azure Synapse Analytics to receive the data.

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:

Actions

Answer Area

Create an external data source for Azure Blob storage.

Create a master key on database.

Enable Transparent Data Encryption.

Create the external table FactSalesOrderDetails.

Load the data to a staging table.

Create an external file format to map the parquet files.

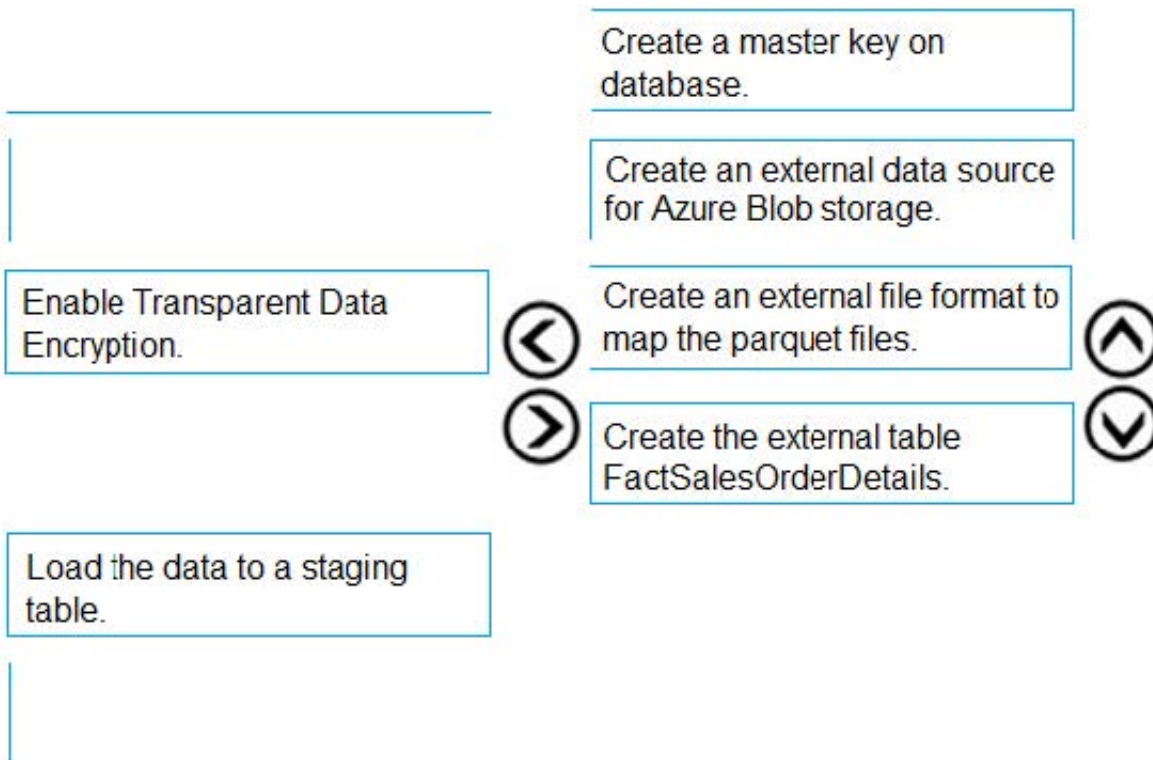


Correct Answer:



Actions

Answer Area



To query the data in your Hadoop data source, you must define an external table to use in Transact-SQL queries. The following steps describe how to configure the external table. Step 1: Create a master key on database.

1.

Create a master key on the database. The master key is required to encrypt the credential secret.

(Create a database scoped credential for Azure blob storage.)

Step 2: Create an external data source for Azure Blob storage.

2.

Create an external data source with CREATE EXTERNAL DATA SOURCE..

Step 3: Create an external file format to map the parquet files.

3.

Create an external file format with CREATE EXTERNAL FILE FORMAT.

Step 4. Create an external table FactSalesOrderDetails

4.

Create an external table pointing to data stored in Azure storage with CREATE EXTERNAL TABLE.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure-blob-storage>

**QUESTION 2**

You have an Azure virtual machine based on a custom image named VM1.

VM1 hosts an instance of Microsoft SQL Server 2019 Standard.

You need to automate the maintenance of VM1 to meet the following requirements:

1.

Automate the patching of SQL Server and Windows Server.

2.

Automate full database backups and transaction log backups of the databases on VM1.

3.

Minimize administrative effort. What should you do first?

A. Enable a system-assigned managed identity for VM1

B. Register VM1 to the Microsoft.Sql resource provider

C. Install an Azure virtual machine Desired State Configuration (DSC) extension on VM1

D. Register VM1 to the Microsoft.SqlVirtualMachine resource provider

Correct Answer: B

Automated Patching depends on the SQL Server infrastructure as a service (IaaS) Agent Extension. The SQL Server IaaS Agent Extension (SqlIaaSExtension) runs on Azure virtual machines to automate administration tasks. The SQL Server

IaaS extension is installed when you register your SQL Server VM with the SQL Server VM resource provider.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-iaas-agent-extension-automate-management>

QUESTION 3

You have an Azure SQL database named DB1.

You need to query the fragmentation information of data and indexes for the tables in DB1.

Which command should you run?

A. `sys.dm_db_index_usage_stats`

B. `DBCC CHECKALLOC`



C. DBCC SHOWCONTIG

D. sts.dm_db_index_physical_stats

Correct Answer: D

sys.dm_db_index_physical_stats (Transact-SQL)

Returns size and fragmentation information for the data and indexes of the specified table or view in SQL Server. For an index, one row is returned for each level of the B-tree in each partition. For a heap, one row is returned for the

IN_ROW_DATA allocation unit of each partition. For large object (LOB) data, one row is returned for the LOB_DATA allocation unit of each partition. If row-overflow data exists in the table, one row is returned for the

ROW_OVERFLOW_DATA allocation unit in each partition.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-db-index-physical-stats-transact-sql>

QUESTION 4

You plan to move two 100-GB databases to Azure.

You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.

What should you use?

- A. two Azure SQL Databases in an elastic pool
- B. two databases hosted in SQL Server on an Azure virtual machine
- C. two databases in an Azure SQL Managed instance
- D. two single Azure SQL databases

Correct Answer: D

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview>

QUESTION 5

DRAG DROP

You have an Azure subscription.

You plan to deploy a new Azure virtual machine that will host a Microsoft SQL Server instance.



You need to configure the disks on the virtual machine. The solution must meet the following requirements:

1.

Minimize latency for transaction logs.

2.

Minimize the impact on IO Of the virtual machine.

Which type of disk should you use for each workload?

To answer, drag the appropriate disk types to the correct workloads. Each disk type may be used once, at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

The screenshot shows a drag-and-drop interface. On the left, under the heading "Disk Types", there are five rectangular buttons: "Local", "Premium SSD", "Standard HDD", "Standard SSD", and "Ultra Disk". To the right of these buttons is a vertical line with five small grey circles. On the right side, under the heading "Answer Area", there are two rows. The first row is labeled "TempDB:" followed by a dashed rectangular box containing the text "Disk Type". The second row is labeled "Transaction logs:" followed by a dashed rectangular box containing the text "Disk Type".

Correct Answer:

**Disk Types**

Local

Standard HDD

Standard SSD

Answer Area

TempDB:

Premium SSD

Transaction logs:

Ultra Disk

Box 1: Premium SSD

Storage: Performance best practices for SQL Server on Azure VMs.

Place tempdb on the local ephemeral SSD (default D:\) drive for most SQL Server workloads that are not part of Failover Cluster Instance (FCI) after choosing the optimal VM size.

Select Premium SSD instead of Standard SSD for better performance.

Box 2: Ultra Disk

For the log drive plan for capacity and test performance versus cost while evaluating the premium P30 - P80 disks

If submillisecond storage latency is required, use Azure ultra disks for the transaction log.

For M-series virtual machine deployments consider write accelerator over using Azure ultra disks.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/performance-guidelines-best-practices-storage>

QUESTION 6

You have an Azure subscription that contain an Azure SQL managed instance named SQLMI1 and a Log Analytics workspace named Workspace1.

You need to collect performance metrics for SQLMI1 and stream the metrics to Workspace1.

- A. Create the private endpoint connection on SQLMI1.
- B. Configure Azure SQL Analytics to use Workspace1.



C. Modify the Computer + storage settings for SQLMI1.

D. Modify the diagnostic settings for SQLMI1.

Correct Answer: D

Instances in Azure SQL Managed Instance

You can set up a managed instance resource to collect the following diagnostic telemetry:

ResourceUsageStats contains vCores count, average CPU percentage, IO requests, bytes read/written, reserved storage space, and used storage space.

To configure streaming of diagnostic telemetry for managed instance and instance databases, you will need to separately configure each:

Enable streaming of diagnostic telemetry for managed instance

Enable streaming of diagnostic telemetry for each instance database

The managed instance container has its own telemetry separate from each instance database's telemetry.

To enable streaming of diagnostic telemetry for a managed instance resource, follow these steps:

1.

Go to the managed instance resource in Azure portal.

2.

Select Diagnostics settings.

3.

Select Turn on diagnostics if no previous settings exist, or select Edit setting to edit a previous setting.

4.

Enter a setting name for your own reference.

5.

Select a destination resource for the streaming diagnostics data: Archive to storage account, Stream to an event hub, or Send to Log Analytics.

6.

For log analytics, select Configure and create a new workspace by selecting +Create New Workspace, or use an existing workspace.

7.

Select the check box for instance diagnostic telemetry: ResourceUsageStats.

8.

Select Save.

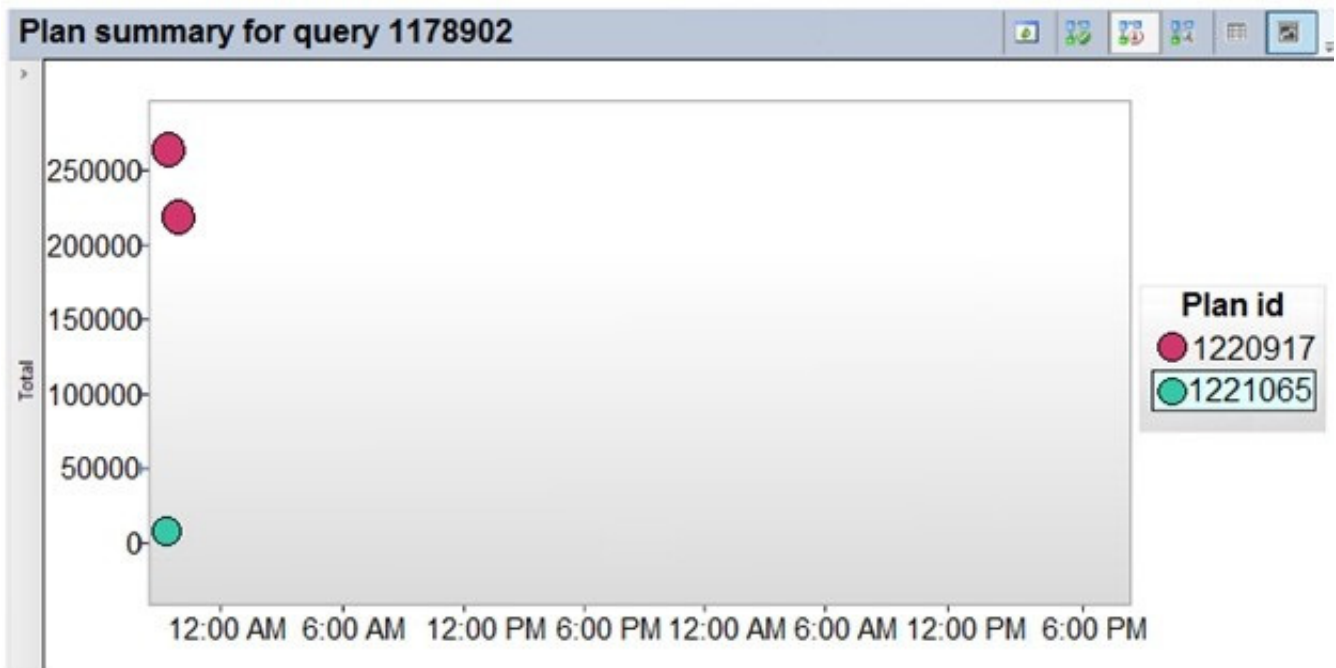


Reference: <https://learn.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-export-configure>

QUESTION 7

You have SQL Server on an Azure virtual machine that contains a database named DB1.

You view a plan summary that shows the duration in milliseconds of each execution of query 1178902 as shown in the following exhibit:



What should you do to ensure that the query uses the execution plan which executes in the least amount of time?

- A. Force the query execution plan for plan 1221065.
- B. Run the DBCC FREEPROCCACHE command.
- C. Force the query execution plan for plan 1220917.
- D. Disable parameter sniffing.

Correct Answer: A

As per exhibit, the execution plan 1221065 has lower execution time compared to plan 1220917.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-store-usage-scenarios>

QUESTION 8

You need to identify the cause of the performance issues on SalesSQLDb1.



Which two dynamic management views should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. sys.dm_pdw_nodes_tran_locks
- B. sys.dm_exec_compute_node_errors
- C. sys.dm_exec_requests
- D. sys.dm_cdc_errors
- E. sys.dm_pdw_nodes_os_wait_stats
- F. sys.dm_tran_locks

Correct Answer: AE

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

A: Use sys.dm_pdw_nodes_tran_locks instead of sys.dm_tran_locks from Azure Synapse Analytics (SQL Data Warehouse) or Parallel Data Warehouse.

E: Example:

The following query will show blocking information.

```
SELECT
t1.resource_type,
t1.resource_database_id,
t1.resource_associated_entity_id,
t1.request_mode,
t1.request_session_id,
t2.blocking_session_id
FROM sys.dm_tran_locks as t1
INNER JOIN sys.dm_os_waiting_tasks as t2
ON t1.lock_owner_address = t2.resource_address;
```

Note: Depending on the system you're working with you can access these wait statistics from one of three locations:

sys.dm_os_wait_stats: for SQL Server sys.dm_db_wait_stats: for Azure SQL Database
sys.dm_pdw_nodes_os_wait_stats: for Azure SQL Data Warehouse

Incorrect Answers:

F: sys.dm_tran_locks returns information about currently active lock manager resources in SQL Server 2019 (15.x). Each row represents a currently active request to the lock manager for a lock that has been granted or is waiting to be



granted.

Instead use `sys.dm_pdw_nodes_tran_locks`.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-tran-locks-transact-sql>

QUESTION 9

Your company uses Azure Stream Analytics to monitor devices.

The company plans to double the number of devices that are monitored.

You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load.

Which metric should you monitor?

- A. Input Deserialization Errors
- B. Late Input Events
- C. Early Input Events
- D. Watermark delay

Correct Answer: D

The Watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far. The watermark delay metric can rise due to:

1.

Not enough processing resources in Stream Analytics to handle the volume of input events.

2.

Not enough throughput within the input event brokers, so they are throttled.

3.

Output sinks are not provisioned with enough capacity, so they are throttled.

Reference: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-time-handling>

QUESTION 10

HOTSPOT

You have an Azure SQL database named DB1 that contains two tables named Table1 and Table2. Both tables contain a column named a Column1. Column1 is used for joins by an application named App1.

You need to protect the contents of Column1 at rest, in transit, and in use.



How should you protect the contents of Column1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Encryption key:

	▼
Column encryption key	
Database encryption key	
Service master key	

Encryption type:

	▼
Deterministic	
Randomized	
Transparent Data Encryption (TDE)	

Correct Answer:

Answer Area

Encryption key:

	▼
Column encryption key	
Database encryption key	
Service master key	

Encryption type:

	▼
Deterministic	
Randomized	
Transparent Data Encryption (TDE)	



Explanation:

Box 1: Column encryption Key

Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.

Incorrect Answers:

TDE encrypts the storage of an entire database by using a symmetric key called the Database Encryption Key (DEK).

Box 2: Deterministic

Always Encrypted is a feature designed to protect sensitive data, such as credit card numbers or national identification numbers (for example, U.S. social security numbers), stored in Azure SQL Database or SQL Server databases. Always

Encrypted allows clients to encrypt sensitive data inside client applications and never reveal the encryption keys to the Database Engine (SQL Database or SQL Server).

Always Encrypted supports two types of encryption: randomized encryption and deterministic encryption.

Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns.

Incorrect Answers:

1.

Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.

2.

Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest. It performs real-time encryption and decryption of the database, associated backups, and transaction log files at rest without requiring changes to the application.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

QUESTION 11

DRAG DROP

You have a database named db1.

The log for db1 contains the following entry.



Date 10/5/2021 10:57:08 AM
Log SQL Server (Current - 10/5/2021 11:26:00 AM)

Sources spid1595

Message

The transaction log for database 'db1' is full due to 'AVAILABILITY_REPLICA'

You need to ensure that db1 can process transactions.

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.

Select and Place:

Actions	Answer Area
Back up the transaction log file.	1 <input type="text"/>
Remove db1 from the availability group.	2 <input type="text"/>
Shrink the transaction log file.	3 <input type="text"/>
Shrink db1.	
Add db1 back to the availability group.	

Correct Answer:

Actions	Answer Area
Back up the transaction log file.	1 Remove db1 from the availability group.
<input type="text"/>	2 Shrink the transaction log file.
<input type="text"/>	3 Add db1 back to the availability group.
Shrink db1.	
<input type="text"/>	



Step 1: Remove db1 from the availability group.

Step 2: Shrink the transaction log file.

You can use SQL Server Management Studio or Transact-SQL to shrink a data or log file.

Example:

SQL

USE UserDB;

GO

DBCC SHRINKFILE (DataFile1, 7);

GO

Incorrect:

* Back up the transaction log file.

Backup the Log file to free up space within the file (logically/virtually).

Step 3: Add db1 back to the availability group.

Note:

Symptoms

Consider the following scenario:

You have Microsoft SQL Server 2012 or a later version installed on a server.

The instance of SQL Server is a primary replica in Always On Availability Groups environment.

The autogrow option for transaction log files is set in SQL Server.

In this scenario, the transaction log may become large and run out of disk space or exceed the MaxSize option set for the transaction log at the primary replica and you receive an error message that resembles the following:

Error: 9002, Severity: 17, State: 9. The transaction log for database \\'%.*Is\\' is full due to \\'AVAILABILITY_REPLICA\\'

Cause

This occurs when the logged changes at primary replica are not yet hardened on the secondary replica.

Workaround

After you identify the secondary database that makes this occur, try one or more of the following methods to work around this issue temporarily:

*

Take the database out of the availability group for the offending secondary.

Note



This method will result in the loss of the High Availability/Disaster Recovery scenario for the secondary. You may have to set up the Availability Group again in the future.

*

If the redo thread is frequently blocked, disable the Readable Secondary feature by changing the ALLOW_CONNECTIONS parameter of the SECONDARY_ROLE for the replica to NO.

Note

This will prevent users from reading the data in the secondary replica which is the root cause of the blocking. Once the redo queue has dropped to an acceptable size, consider enabling the feature again.

*

Enable the autogrow setting if it is disabled and there is available disk space.

*

Increase the MaxSize value for the transaction log file if it has been reached and there is available disk space.

*

Add an additional transaction log file if the current one has reached the system maximum of 2 TB or if additional space is available on another available volume.

Reference: <https://learn.microsoft.com/en-us/troubleshoot/sql/availability-groups/error-9002-transaction-log-large>
<https://thecloudtechnologist.com/2015/03/18/shrinking-sql-log-files-in-an-availability-group-cluster-or-database-mirror/>

QUESTION 12

HOTSPOT

You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD service principal named app1.

You need to ensure that app1 can access db1. The solution must use the principle of least privilege.

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

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

CREATE

	▼	[app1]
CREDENTIAL		
LOGIN		
USER		

	▼
FOR LOGIN app1	
FROM EXTERNAL PROVIDER	
FROM LOGIN app1	
WITHOUT LOGIN	

Correct Answer:



Answer Area

CREATE

	▼	[app1]
CREDENTIAL		
LOGIN		
USER		

	▼	
FOR LOGIN app1		
FROM EXTERNAL PROVIDER		
FROM LOGIN app1		
WITHOUT LOGIN		

Box 1: USER

Azure Active Directory service principal with Azure SQL

When an Azure AD application is registered using the Azure portal or a PowerShell command, two objects are created in the Azure AD tenant:

An application object

A service principal object

SQL Database and SQL Managed Instance support the following Azure AD objects:

Azure AD users (managed, federated, and guest)

Azure AD groups (managed and federated)

Azure AD applications

The T-SQL command CREATE USER [Azure_AD_Object] FROM EXTERNAL PROVIDER on behalf of an Azure AD application is now supported for SQL Database.

Box 2: FROM EXTERNAL PROVIDER

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal>

**QUESTION 13**

You are designing an enterprise data warehouse in Azure Synapse Analytics that will contain a table named Customers. Customers will contain credit card information.

You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers. The solution must prevent all the salespeople from viewing or inferring the credit card information.

What should you include in the recommendation?

- A. row-level security
- B. data masking
- C. Always Encrypted
- D. column-level security

Correct Answer: D

inferring work is the key. With data-masking you cannot prevent to ask if a column is equal to a certain value. Also, the salespeople don't need any information of the credit card, maybe customer support would have a use of the masked credit card. So, column-level security is the solution.

QUESTION 14

You have an Azure subscription that contains two instances of SQL Server on Azure Virtual Machines named VM1 and VM2. Both instances run Microsoft SQL Server 2019 CU8.

You need to deploy a failover cluster instance (FCI) to VM1 and VM2 that will use Azure shared disks. The solution must maximize resiliency.

Which quorum option should you use?

- A. node majority with a cloud witness
- B. node majority with no witness
- C. node majority with a file share witness
- D. node majority with a disk witness

Correct Answer: D

Configure quorum

Since the disk witness is the most resilient quorum option, and the FCI solution uses Azure shared disks, it's recommended to configure a disk witness as the quorum solution.

Reference: <https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/failover-cluster-instance-azure-shared-disks-manually-configure>



QUESTION 15

You plan to deploy an app that includes an Azure SQL database and an Azure web app. The app has the following requirements:

1.

The web app must be hosted on an Azure virtual network.

2.

The Azure SQL database must be assigned a private IP address.

3.

The Azure SQL database must allow connections only from the virtual network.

You need to recommend a solution that meets the requirements.

What should you include in the recommendation?

A. Azure Private Link

B. a network security group (NSG)

C. a database-level firewall

D. a server-level firewall

Correct Answer: A

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/private-endpoint-overview>

[DP-300 Practice Test](#)

[DP-300 Exam Questions](#)

[DP-300 Braindumps](#)