



DUMPSHQ

Free Questions for DP-300 by dumpshq

Shared by Ortega on 05-09-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

D18912E1457D5D1DDCBD40AB3BF70D5D

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.

Options:

- 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

Answer:

A, E

Explanation:

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.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

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

Question 2

Question Type: Hotspot

You need to recommend a configuration for ManufacturingSQLDb1 after the migration to Azure. The solution must meet the business requirements.

What should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Quorum model:

	▼
Cloud witness	
Disk witness	
File share witness	

Azure resource for the availability group listener:

	▼
Azure Application Gateway	
Azure Basic Load Balancer	

Scenario: Business Requirements

Litware identifies business requirements include: meet an SLA of 99.99% availability for all Azure deployments.

Box 1: Cloud witness

If you have a Failover Cluster deployment, where all nodes can reach the internet (by extension of Azure), it is recommended that you configure a Cloud Witness as your quorum witness resource.

Box 2: Azure Basic Load Balancer

Microsoft guarantees that a Load Balanced Endpoint using Azure Standard Load Balancer, serving two or more Healthy Virtual Machine Instances, will be available 99.99% of the time.

Note: There are two main options for setting up your listener: external (public) or internal. The external (public) listener uses an internet facing load balancer and is associated with a public Virtual IP (VIP) that is accessible over the internet. An internal listener uses an internal load balancer and only supports clients within the same Virtual Network.

Answer:

Explanation:

<https://technet.microsoft.com/windows-server-docs/failover-clustering/deploy-cloud-witness>

https://azure.microsoft.com/en-us/support/legal/sla/load-balancer/v1_0/

Question 3

Question Type: MultipleChoice

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 have an Azure SQL database named Sales.

You need to implement disaster recovery for Sales to meet the following requirements:

- * During normal operations, provide at least two readable copies of Sales.
- * Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and failover groups.

Does this meet the goal?

Options:

A) Yes

B) No

Answer:

B

Explanation:

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

<https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

Question 4

Question Type: MultipleChoice

D18912E1457D5D1DDCBD40AB3BF70D5D

You need to implement authentication for ResearchDB1. The solution must meet the security and compliance requirements.

What should you run as part of the implementation?

Options:

- A) CREATE LOGIN and the FROM WINDOWS clause
- B) CREATE USER and the FROM CERTIFICATE clause
- C) CREATE USER and the FROM LOGIN clause
- D) CREATE USER and the ASYMMETRIC KEY clause
- E) CREATE USER and the FROM EXTERNAL PROVIDER clause

Answer:

E

Explanation:

Scenario: Authenticate database users by using Active Directory credentials.

(Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.)

Authenticate the user in SQL Database or SQL Data Warehouse based on an Azure Active Directory user:

```
CREATE USER [Fritz@contoso.com] FROM EXTERNAL PROVIDER;
```

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql>

Question 5

Question Type: MultipleChoice

D18912E1457D5D1DDCBD40AB3BF70D5D

What should you do after a failover of SalesSQLDb1 to ensure that the database remains accessible to SalesSQLDb1App1?

Options:

- A) Configure SalesSQLDb1 as writable.
- B) Update the connection strings of SalesSQLDb1App1.
- C) Update the firewall rules of SalesSQLDb1.
- D) Update the users in SalesSQLDb1.

Answer:

C

Explanation:

Scenario: SalesSQLDb1 uses database firewall rules and contained database users.

Question 6

Question Type: MultipleChoice

D18912E1457D5D1DDCBD40AB3BF70D5D

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.

Options:

- 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

Answer:

A, E

Explanation:

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.

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

Question 7

Question Type: MultipleChoice

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 have an Azure SQL database named Sales.

You need to implement disaster recovery for Sales to meet the following requirements:

* During normal operations, provide at least two readable copies of Sales.

* Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and geo-replication.

Does this meet the goal?

Options:

A) Yes

B) No

Answer:

B

Explanation:

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

<https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

To Get Premium Files for DP-300 Visit

<https://www.p2pexams.com/products/dp-300>

For More Free Questions Visit

<https://www.p2pexams.com/microsoft/pdf/dp-300>

