

**100%** Money Back  
**Guarantee**

**Vendor:**Microsoft

**Exam Code:**DP-300

**Exam Name:**Administering Relational Databases on  
Microsoft Azure

**Version:**Demo

## QUESTION 1

You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2.

You need to limit the number of IOPs that App2 queries generate on SQL1.

Which two actions should you perform on SQL1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable query optimizer fixes.
- B. Enable Resource Governor.
- C. Enable parameter sniffing.
- D. Create a workload group.
- E. Configure In-memory OLTP.
- F. Run the Database Engine Tuning Advisor.
- G. Reduce the Max Degree of Parallelism value.

Correct Answer: BC

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server-ver15>

---

## QUESTION 2

You have an Azure SQL Database server named sqlsrv1 that hosts 10 Azure SQL databases.

The databases perform slower than expected.

You need to identify whether the performance issue relates to the use of tempdb by Azure SQL databases on sqlsrv1.

What should you do?

- A. Run Query Store-based queries
- B. Review information provided by SQL Server Profiler-based traces
- C. Review information provided by Query Performance Insight
- D. Run dynamic management view-based queries

Correct Answer: D

Microsoft Azure SQL Database and Azure SQL Managed Instance enable a subset of dynamic management views to diagnose performance problems, which might be caused by blocked or long-running queries, resource bottlenecks,

poor

query plans, and so on.

This include edentifying tempdb performance issues:

When identifying IO performance issues, the top wait types associated with tempdb issues is PAGELATCH\_\* (not PAGEIOLATCH\_\*).

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs>

---

### QUESTION 3

You have an Azure subscription that contains three instances of SQL Server on Azure Virtual Machines.

You plan to implement a disaster recovery solution.

You need to be able to perform disaster recovery drills regularly. The solution must meet the following requirements:

1.

Minimize administrative effort for the recovery drills.

2.

Isolate the recovery environment from the production environment What should you use?

- A. native Microsoft SQL Server backup
- B. Azure Site Recovery
- C. Recovery Services vaults
- D. Azure Backup

Correct Answer: B

Set up disaster recovery for SQL Server

You can protect the SQL Server back end of an application. You do so by using a combination of SQL Server business continuity and disaster recovery (BCDR) technologies and Azure Site Recovery.

SQL Server disaster recovery capabilities include:

Failover clustering

Always On availability groups

Database mirroring

Log shipping

Active geo-replication

Auto-failover groups

Note: Azure Recovery Services contributes to your BCDR strategy:

Site Recovery service: Site Recovery helps ensure business continuity by keeping business apps and workloads running during outages. Site Recovery replicates workloads running on physical and virtual machines (VMs) from a primary site

to a secondary location. When an outage occurs at your primary site, you fail over to a secondary location, and access apps from there. After the primary location is running again, you can fail back to it.

Backup service: The Azure Backup service keeps your data safe and recoverable.

Site Recovery can manage replication for:

Azure VMs replicating between Azure regions  
Replication from Azure Public Multi-Access Edge Compute (MEC) to the region  
Replication between two Azure Public MECs  
On-premises VMs, Azure Stack VMs, and physical servers

Reference: <https://learn.microsoft.com/en-us/azure/site-recovery/site-recovery-sql> <https://learn.microsoft.com/en-us/azure/site-recovery/site-recovery-overview>

---

#### QUESTION 4

You have an Azure SQL Database managed instance.

The instance starts experiencing performance issues.

You need to identify which query is causing the issue and retrieve the execution plan for the query. The solution must minimize administrative effort.

What should you use?

- A. the Azure portal
- B. Extended Events
- C. Query Store
- D. dynamic management views

Correct Answer: D

Use the dynamic management view `sys.dm_exec_requests` to track currently executing queries and the associated worker time.

Incorrect:

Not C: DMVs that track Query Store and wait statistics show results for only successfully completed and timed-out queries.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/identify-query-performance-issues>

---

## QUESTION 5

Which counter should you monitor for real-time processing to meet the technical requirements?

- A. SU% Utilization
- B. CPU% utilization
- C. Concurrent users
- D. Data Conversion Errors

Correct Answer: B

Scenario: Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.

To monitor the performance of a database in Azure SQL Database and Azure SQL Managed Instance, start by monitoring the CPU and IO resources used by your workload relative to the level of database performance you chose in selecting a particular service tier and performance level.

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

---

## QUESTION 6

You have 50 Azure SQL databases.

You need to notify the database owner when the database settings, such as the database size and pricing tier, are modified in Azure.

What should you do?

- A. Create a diagnostic setting for the activity log that has the Security log enabled.
- B. For the database, create a diagnostic setting that has the InstanceAndAppAdvanced metric enabled.
- C. Create an alert rule that uses a Metric signal type.
- D. Create an alert rule that uses an Activity Log signal type.

Correct Answer: D

Activity log events - An alert can trigger on every event, or, only when a certain number of events occur. Incorrect Answers:

C: Metric values - The alert triggers when the value of a specified metric crosses a threshold you assign in either direction. That is, it triggers both when the condition is first met and then afterwards when that condition is no longer being met.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/alerts-insights-configure-portal>

---

## QUESTION 7

You are designing a star schema for a dataset that contains records of online orders. Each record includes an order date, an order due date, and an order ship date.

You need to ensure that the design provides the fastest query times of the records when querying for arbitrary date ranges and aggregating by fiscal calendar attributes.

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

NOTE: Each correct selection is worth one point.

- A. Create a date dimension table that has a DateTime key.
- B. Create a date dimension table that has an integer key in the format of YYYYMMDD.
- C. Use built-in SQL functions to extract date attributes.
- D. Use integer columns for the date fields.
- E. Use DateTime columns for the date fields.

Correct Answer: BD

Reference:

[https://community.idera.com/database-tools/blog/b/community\\_blog/posts/why-use-a-date-dimension-table-in-a-data-warehouse](https://community.idera.com/database-tools/blog/b/community_blog/posts/why-use-a-date-dimension-table-in-a-data-warehouse)

---

## QUESTION 8

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 9**

**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.

Hot Area:

**Answer Area**

Quorum mode:

	▼
Node majority with witness	
Node majority	
No majority	

Azure resource for the availability group listener:

	▼
Azure Application Gateway	
Azure Basic Load Balancer	
Azure Standard Load Balancer	

Correct Answer:

## Answer Area

Quorum mode:

	▼
Node majority with witness	
Node majority	
No majority	

Azure resource for the availability group listener:

	▼
Azure Application Gateway	
Azure Basic Load Balancer	
Azure Standard Load Balancer	

Box 1: Node majority with witness

As a general rule when you configure a quorum, the voting elements in the cluster should be an odd number. Therefore, if the cluster contains an even number of voting nodes, you should configure a disk witness or a file share witness.

Note: Mode: Node majority with witness (disk or file share)

Nodes have votes. In addition, a quorum witness has a vote. The cluster quorum is the majority of voting nodes in the active cluster membership plus a witness vote. A quorum witness can be a designated disk witness or a designated file share witness.

Box 2: Azure Standard 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.

Scenario: Business Requirements

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

Incorrect Answers:

Basic Balancer: No SLA is provided for Basic Load Balancer.

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.

Reference:

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



---

## QUESTION 10

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

You have an application that queries DB1 to generate a sales report.

You need to see the parameter values from the last time the query was executed.

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

NOTE: Each correct selection is worth one point.

- A. Enable Last\_Query\_Plan\_Stats in the master database
- B. Enable Lightweight\_Query\_Profiling in DB1
- C. Enable Last\_Query\_Plan\_Stats in DB1
- D. Enable Lightweight\_Query\_Profiling in the master database
- E. Enable PARAMETER\_SNIFFING in DB1

Correct Answer: BC

Last\_Query\_Plan\_Stats allows you to enable or disable collection of the last query plan statistics (equivalent to an actual execution plan) in sys.dm\_exec\_query\_plan\_stats.

Lightweight profiling can be disabled at the database level using the LIGHTWEIGHT\_QUERY\_PROFILING database scoped configuration: ALTER DATABASE SCOPED CONFIGURATION SET LIGHTWEIGHT\_QUERY\_PROFILING = OFF;.

Incorrect Answers:

A: Enable it for DB1, not for the master database.

E: Parameter sensitivity, also known as "parameter sniffing", refers to a process whereby SQL Server "sniffs" the current parameter values during compilation or recompilation, and passes it along to the Query Optimizer so that they can be used to generate potentially more efficient query execution plans.

Parameter values are sniffed during compilation or recompilation for the following types of batches: Stored procedures  
Queries submitted via sp\_executesql Prepared queries

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure>  
<https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure>

---

## QUESTION 11

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 12

You create five Azure SQL Database instances on the same logical server.

In each database, you create a user for an Azure Active Directory (Azure AD) user named User1.

User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.

You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.

What should you do?

- A. Create User1 in the master database.
- B. Assign User1 the db\_datareader role for the master database.
- C. Assign User1 the db\_datareader role for the databases that User1 creates.
- D. Grant SELECT on sys.databases to public in the master database.

Correct Answer: A

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/logins-create-manage>