

[Download Full Version DP-300 Exam Dumps\(Updated in Feb/2023\)](#)

QUESTION 50

You have an Azure SQL database named db1 on a server named server1.

The Intelligent Insights diagnostics log identifies that several tables are missing indexes.

You need to ensure that indexes are created for the tables.

What should you do?

- A. Run the DBCC SQLPERF command.
- B. Run the dbcc dbreindex command.
- C. Modify the automatic tuning settings for db1.
- D. Modify the Query Store settings for db1.

Correct Answer: C

Explanation:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-overview>

QUESTION 51

You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1.

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1. The on-premises network connects to VNET1 by using an ExpressRoute connection.

You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service.

You need to configure VNET1 to support the migration.

What should you do?

- A. Configure service endpoints.
- B. Configure virtual network peering.
- C. Deploy an Azure firewall.
- D. Configure network security groups (NSGs).

Correct Answer: A

Explanation:

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-to-managed-instance>

QUESTION 52

You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.

You scale up the virtual machine to 16 vCPUs and 64 GB of memory.

You need to provide the lowest latency for tempdb.

What is the total number of data files that tempdb should contain?

- A. 2
- B. 4

[DP-300 Exam Dumps](#) [DP-300 PDF Dumps](#) [DP-300 VCE Dumps](#) [DP-300 Q&As](#)

<https://www.ensurepass.com/DP-300.html>

[Download Full Version DP-300 Exam Dumps\(Updated in Feb/2023\)](#)

- C. 8
- D. 64

Correct Answer: D

Explanation:

The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database>

QUESTION 53

HOTSPOT

You have a 50-TB Microsoft SQL Server database named DB1.

You need to reduce the time it takes to perform database consistency checks of DB1.

Which Transact-SQL command should you run? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

DBCC CHECKDB ([DB1],

▼
NOINDEX
REPAIR_FAST
REPAIR_REBUILD

) with

▼
ALL_ERRORMSGs
NO_INFOMSGs
PHYSICAL_ONLY

Correct Answer:

DBCC CHECKDB ([DB1],

▼
NOINDEX
REPAIR_FAST
REPAIR_REBUILD

) with

▼
ALL_ERRORMSGs
NO_INFOMSGs
PHYSICAL_ONLY

[Download Full Version DP-300 Exam Dumps\(Updated in Feb/2023\)](#)

QUESTION 54

DRAG DROP

You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1.

You plan to deploy an Azure web app named webapp1 that will export rows in Table1 that have changed.

You need to ensure that webapp1 can identify the changes to Table1. The solution must meet the following requirements:

- Minimize compute times.
- Minimize storage.

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.

Actions

- From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGETABLE()` function.
- Connect to SQLDb1 and run the following Transact-SQL statement.
`ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON`
- From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGE_TRACKING_CURRENT_VERSION()` function.
- Connect to SQLDb1 and run the following Transact-SQL statement.
`EXEC sys.sp_cdc_enable_table`
- Connect to SQLDb1 and run the following Transact-SQL statement.
`EXEC sys.sp_cdc_enable_db`
- Connect to SQLDb1 and run the following Transact-SQL statement.
`ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING`

Answer Area

Navigation: > <

Correct Answer:

Actions

- From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGETABLE()` function.
- Connect to SQLDb1 and run the following Transact-SQL statement.
`ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON`
- From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGE_TRACKING_CURRENT_VERSION()` function.
- Connect to SQLDb1 and run the following Transact-SQL statement.
`EXEC sys.sp_cdc_enable_table`
- Connect to SQLDb1 and run the following Transact-SQL statement.
`EXEC sys.sp_cdc_enable_db`
- Connect to SQLDb1 and run the following Transact-SQL statement.
`ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING`

Answer Area

Navigation: > <

Sequence of actions in Answer Area:

- Connect to SQLDb1 and run the following Transact-SQL statement.
`ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON`
- Connect to SQLDb1 and run the following Transact-SQL statement.
`ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING`
- From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGE` function.

[Download Full Version DP-300 Exam Dumps\(Updated in Feb/2023\)](#)

QUESTION 55

HOTSPOT

You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

Configure

Feedback

Compute Hardware

Click "Change configuration" to see details for all hardware generations available including memory optimized and compute optimized options

Hardware Configuration

Gen5
up to 40 vCores, up to 120 GB memory
[Change configuration](#)

Max vCores

Min vCores

Auto-pause delay

The database automatically pauses if it is inactive for the time period specified here, and automatically resumes when database activity recurs. Alternatively, auto-pausing can be disabled.

☐ Enable auto-pause

Days Hours Minutes

Data max size

Apply

Cost summary

Gen5 - General Purpose (GP_5, Gen5_0)	
Cost per GB (in USD)	0.12
Max storage selected (in GB)	x 1040
ESTIMATED STORAGE COST / MONTH	119.60 USD
COMPUTE COST / VCORE / SECOND¹	0.000145 USD

NOTES

¹ Serverless databases are billed in vCores based on a combination of CPU and memory utilization. [Learn more about serverless billing](#)

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

After four hours of inactivity, the database requires [answer choice] to resume operations for new activities.

- no extra time
- up to 10 minutes
- up to one minute

The database configuration reduces the cost of **[answer choice]** usage patterns.

intermittent and unpredictable
regular and high
steady and low

Correct Answer:

[Download Full Version DP-300 Exam Dumps\(Updated in Feb/2023\)](#)

After four hours of inactivity, the database requires [answer choice] to resume operations for new activities.

no extra time
up to 10 minutes
up to one minute

The database configuration reduces the cost of [answer choice] usage patterns.

intermittent and unpredictable
regular and high
steady and low

QUESTION 56

You have an Azure subscription that uses a domain named contoso.com.

You have two Azure VMs named DBServer1 and DBServer2. Each of them hosts a default SQL Server instance. DBServer1 is in the East US Azure region and contains a database named DatabaseA. DBServer2 is in the West US Azure region.

DBServer1 has a high volume of data changes and low latency requirements for data writes.

You need to configure a new availability group for DatabaseA. The secondary replica will reside on DBServer2.

What should you do?

- A. Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Asynchronous.
- B. Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Synchronous.
- C. Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Asynchronous.
- D. Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Synchronous.

Correct Answer: C

Explanation:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on-availability-groups?view=sql-server-ver15>

QUESTION 57

HOTSPOT

You have an Azure Data Lake Storage Gen2 account named account1 that stores logs as shown in the following table.

Type	Designated retention period
Application	360 days
Infrastructure	60 days

You do not expect that the logs will be accessed during the retention periods.

You need to recommend a solution for account1 that meets the following requirements:

Automatically deletes the logs at the end of each retention period

[DP-300 Exam Dumps](#) [DP-300 PDF Dumps](#) [DP-300 VCE Dumps](#) [DP-300 Q&As](#)

<https://www.ensurepass.com/DP-300.html>