



## **Microsoft**

### **Exam Questions DP-300**

Administering Relational Databases on Microsoft Azure (beta)

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

**NEW QUESTION 1**

- (Exam Topic 5)

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
- > Minimizes storage costs

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.

To minimize storage costs:

Store the infrastructure logs and the application logs in the Archive access tier.

Store the infrastructure logs and the application logs in the Cool access tier.

Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

Azure Data Factory pipelines

Azure Blob storage lifecycle management rules

Immutable Azure Blob storage time-based retention policies

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

A picture containing text Description automatically generated

Box 1: Store the infrastructure logs in the Cool access tier the application logs in the Archive access tier Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Box 2: Azure Blob storage lifecycle management rules

Blob storage lifecycle management offers a rich, rule-based policy that you can use to transition your data to the best access tier and to expire data at the end of its lifecycle.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

**NEW QUESTION 2**

- (Exam Topic 5)

You have an Azure SQL database.

You discover that the plan cache is full of compiled plans that were used only once.

You run theselect \* from sys.database\_scoped\_configurationsTransact-SQL command and receive the results shown in the following table.

configuration_id	name	value	is_value_default
1	LEGACY_CARDINALITY_ESTIMATION	0	1
2	QUERY_OPTIMIZER_HOTFIXES	0	1
3	OPTIMIZE_FOR_AD_HOC_WORKLOADS	0	1
4	ACCELERATED_PLAN_FORCING	1	1

You need relieve the memory pressure. What should you configure?

- A. LEGACY\_CARDINALITY\_ESTIMATION
- B. QUERY\_OPTIMIZER\_HOTFIXES
- C. OPTIMIZE\_FOR\_AD\_HOC\_WORKLOADS
- D. ACCELERATED\_PLAN\_FORCING

**Answer:** C

**Explanation:**

OPTIMIZE\_FOR\_AD\_HOC\_WORKLOADS = { ON | OFF }

Enables or disables a compiled plan stub to be stored in cache when a batch is compiled for the first time. The default is OFF. Once the database scoped configuration OPTIMIZE\_FOR\_AD\_HOC\_WORKLOADS is enabled for a database, a compiled plan stub will be stored in cache when a batch is compiled for the first time. Plan stubs have a smaller memory footprint compared to the size of the full compiled plan.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql>

**NEW QUESTION 3**

- (Exam Topic 5)

You configure version control for an Azure Data Factory instance as shown in the following exhibit.

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.

Azure Resource Manager (ARM) templates for the pipeline assets as stored in

- /
- adf\_publish
- main
- Parameterization template

A Data Factory Azure Resource Manager (ARM) template named contososales can be found in

- /contososales
- /dwh\_batchetl/adf\_publish/contososales
- /main

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, text, application Description automatically generated

Box 1: adf\_publish

By default, data factory generates the Resource Manager templates of the published factory and saves them into a branch called adf\_publish. To configure a custom publish branch, add a publish\_config.json file to the root folder in the collaboration branch. When publishing, ADF reads this file, looks for the field publishBranch, and saves all Resource Manager templates to the specified location. If the branch doesn't exist, data factory will automatically create it. An example of what this file looks like is below:

```
{
  "publishBranch": "factory/adf_publish"
}
```

Box 2: /dwh\_barchlet/ adf\_publish/contososales

RepositoryName: Your Azure Repos code repository name. Azure Repos projects contain Git repositories to manage your source code as your project grows. You can create a new repository or use an existing repository that's already in your project.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/source-control>

**NEW QUESTION 4**

- (Exam Topic 5)

You have two on-premises servers that run Windows Server 2019 and host a Microsoft SQL Server 2017 Always On availability group named AG1. AG1 contains a single database named DB1.

You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Linux.

You need to migrate DB1 to a SQL Server 2019 instance on VM1. The solution must minimize the downtime of DB1 during the migration.

What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

To prepare for the migration:

To perform the migration, use:

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**  
**Answer Area**

To prepare for the migration:

To perform the migration, use:

**NEW QUESTION 5**

- (Exam Topic 5)

Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks. You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions in the answer area and arrange them in the correct order.

**Actions**

**Answer Area**

- Deploy Grafana to an Azure virtual machine.
- Build a **spark-listeners-loganalytics-1.0-SNAPSHOT.jar** JAR file.
- Create Dropwizard counters in the application code.
- Create a data source in Azure Monitor.
- Configure the Databricks cluster to use the Databricks monitoring library.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, text, application Description automatically generated with medium confidence

Send application metrics using Dropwizard.

Spark uses a configurable metrics system based on the Dropwizard Metrics Library.

To send application metrics from Azure Databricks application code to Azure Monitor, follow these steps: Step 1: Configure your Azure Databricks cluster to use the Databricksmonitoring library.

Prerequisite: Configure your Azure Databricks cluster to use the monitoring library. Step 2: Build the spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file

Step 3: Create Dropwizard counters in your application code Create Dropwizard gauges or counters in your application code

**NEW QUESTION 6**

- (Exam Topic 5)

You have an Azure subscription that is linked to an Azure AD tenant named contoso.com. The subscription contains an Azure SQL database named SQL 1 and an Azure web named app1. App1 has the managed identity feature enabled. You need to create a new database user for app1.

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.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

<https://learn.microsoft.com/en-us/azure/app-service/tutorial-connect-msi-sql-database?tabs=windowsclient%2Ce>

**NEW QUESTION 7**

- (Exam Topic 5)

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.

**Answer:** BD

**Explanation:**

Reference:

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

**NEW QUESTION 8**

- (Exam Topic 5)

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines. The virtual machine hosts a database named DB1. You need to monitor DB1 by using Extended Events. The solution must meet the following requirements:

- Capture raw event data and store the data in Azure Storage.
- Minimize the performance impact of capturing extended events.

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.

**Answer Area**

```
CREATE EVENT SESSION session1 ON DATABASE
ADD EVENT sqlserver.sql_statement_starting
(
ACTION (sqlserver.sql_text)
WHERE statement LIKE 'UPDATE gmTabEmployee%'
)
ADD TARGET
package0.
event_file
event_file
event_stream
ring_buffer
)
SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'
)
WITH
(MAX_MEMORY = 10 MB,
EVENT_RETENTION_MODE=
ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_SINGLE_EVENT_LOSS
NO_EVENT_LOSS
MAX_DISPATCH_LATENCY = 3 SECONDS)
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**  
**Answer Area**

```
CREATE EVENT SESSION session1 ON DATABASE
ADD EVENT sqlserver.sql_statement_starting
(
ACTION (sqlserver.sql_text)
WHERE statement LIKE 'UPDATE gmTabEmployee%'
)
ADD TARGET
package0.
event_file
event_file
event_stream
ring_buffer
)
SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'
)
WITH
(MAX_MEMORY = 10 MB,
EVENT_RETENTION_MODE=
ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_SINGLE_EVENT_LOSS
NO_EVENT_LOSS
MAX_DISPATCH_LATENCY = 3 SECONDS)
```

**NEW QUESTION 9**

- (Exam Topic 5)

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a user named User1. SQLVM1 hosts a database named DB1.

You need to ensure that User1 can perform the following tasks on DB1:

- Create jobs.
  - View all jobs.
  - Modify, delete, and disable the jobs the user created. The solution must use the principle of least privilege.
- Which built-in database role should you assign to User1, and where is the role defined? To answer, select the appropriate options in the answer area.  
 NOTE: Each correct selection is worth one point.

### Answer Area

Location:

DB1
master
msdb

Built-in role:

SQLAgentOperatorRole
SQLAgentReaderRole
SQLAgentUserRole

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

<https://learn.microsoft.com/en-us/sql/ssms/agent/sql-server-agent-fixed-database-roles?view=sql-server-ver16#s>

**NEW QUESTION 10**

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2016 server named Server1 that contains a database named DB1. You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service. How should you configure the backup of DB1? To answer, select the appropriate options in the answer area. NOTE:Each correct selection is worth one point.

### Answer Area

Backup type:

Full and log backups only
Full backup only
Log backup only

Backup option:

WITH CHECKSUM
WITH NOINIT
WITH UNLOAD

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Full and log backups only

Make sure to take every backup on a separate backup media (backup files). Azure Database Migration Service doesn't support backups that are appended to a single backup file. Take full backup and log backups to separate backup files.

Box 2: WITH CHECKSUM

Azure Database Migration Service uses the backup and restore method to migrate your on-premises databases to SQL Managed Instance. Azure Database Migration Service only supports backups created using checksum.

Reference:

<https://docs.microsoft.com/en-us/azure/dms/known-issues-azure-sql-db-managed-instance-online>

**NEW QUESTION 10**

- (Exam Topic 5)

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?

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

**NEW QUESTION 12**

- (Exam Topic 5)

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 dbreindexcommand.
- C. Modify the automatic tuning settings for db1.
- D. Modify the Query Store settings for db1.

**Answer: C**

**Explanation:**

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

**NEW QUESTION 14**

- (Exam Topic 5)

You have a SQL Server on Azure Virtual Machines instance that hosts a 10-TB SQL database named DB1. You need to identify and repair any physical or logical corruption in DB1. The solution must meet the following requirements:

- Minimize how long it takes to complete the procedure.
- Minimize data loss.

How should you complete the command? To answer, select the appropriate options in the answer area NOTE: Each correct selection is worth one point.

**Answer Area**

DBCC CHECK [DB1], 

NOINDEX
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST
REPAIR_REBUILD

 ) WITH 

EXTENDED_LOGICAL_CHECKS;
PHYSICAL_ONLY;
TABLOCK;

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

**Answer Area**

DBCC CHECK [DB1], 

NOINDEX
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST
REPAIR_REBUILD

 ) WITH 

EXTENDED_LOGICAL_CHECKS;
PHYSICAL_ONLY;
TABLOCK;

**NEW QUESTION 15**

- (Exam Topic 5)

You need to migrate an on-premises Microsoft SQL Server database to Azure SQL Database. The solution must minimize downtime.

What should you do?

- A. Configure Transaction Log Shipping.
- B. Implement Always On availability groups.
- C. Configure transactional replication.
- D. Import a BACPAC.

**Answer: C**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/migrate-to-database-from-sql-server#method-1-migra>

**NEW QUESTION 20**

- (Exam Topic 5)

You have an Azure Data Lake Storage Gen2 container.

Data is ingested into the container, and then transformed by a data integration application. The data is NOT modified after that. Users can read files in the container but cannot modify the files.

You need to design a data archiving solution that meets the following requirements:

- > New data is accessed frequently and must be available as quickly as possible.
- > Data that is older than five years is accessed infrequently but must be available within one second when requested.
- > Data that is older than seven years is NOT accessed. After seven years, the data must be persisted at the lowest cost possible.
- > Costs must be minimized while maintaining the required availability.

How should you manage the data? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Five-year-old data:

Delete the blob.
Move to archive storage.
Move to cool storage.
Move to hot storage.

Seven-year-old data:

Delete the blob.
Move to archive storage.
Move to cool storage.
Move to hot storage.

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Text, table Description automatically generated

Box 1: Move to cool storage

The cool access tier has lower storage costs and higher access costs compared to hot storage. This tier is intended for data that will remain in the cool tier for at least 30 days. Example usage scenarios for the cool access tier include:

Short-term backup and disaster recovery

Older data not used frequently but expected to be available immediately when accessed

Large data sets that need to be stored cost effectively, while more data is being gathered for future processing Note: Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Box 2: Move to archive storage

Example usage scenarios for the archive access tier include: Long-term backup, secondary backup, and archival datasets

Original (raw) data that must be preserved, even after it has been processed into final usable form Compliance and archival data that needs to be stored for a long time and is hardly ever accessed Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

**NEW QUESTION 25**

- (Exam Topic 5)

You have an Azure SQL database named DB1. You run a query while connected to DB1.

You review the actual execution plan for the query, and you add an index to a table referenced by the query. You need to compare the previous actual execution plan for the query to the Live Query Statistics.

What should you do first in Microsoft SQL Server Management Studio (SSMS)?

- A. For DB1, set QUERY\_CAPTURE\_MODE of Query Store to All.
- B. Run the SET SHOWPLAN\_ALL Transact-SQL statement.
- C. Save the actual execution plan.
- D. Enable Query Store for DB1.

**Answer: C**

**Explanation:**

The Plan Comparison menu option allows side-by-side comparison of two different execution plans, for easier identification of similarities and changes that explain the different behaviors for all the reasons stated above.

This option can compare between:

Two previously saved execution plan files (.sqlplan extension).

One active execution plan and one previously saved query execution plan. Two selected query plans in Query Store.

**NEW QUESTION 27**

- (Exam Topic 5)

You have an on-premises datacenter that contains a 14-TB Microsoft SQL Server database.

You plan to create an Azure SQL managed instance and migrate the on-premises database to the new instance. Which three service tiers support the SQL managed instance? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. General Purpose Standard
- B. Business Critical Premium
- C. Business Critical Memory Optimized Premium
- D. General Purpose Premium
- E. Business Critical Standard

**Answer:** BCD

**NEW QUESTION 31**

- (Exam Topic 5)

You plan to migrate on-premises Microsoft SQL Server databases to Azure.

You need to identify which deployment and resiliency options meet the following requirements:

- > Support user-initiated backups.
- > Support multiple automatically replicated instances across Azure regions.
- > Minimize administrative effort to implement and maintain business continuity. What should you identify? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Deployment option:

	▼
Azure SQL Managed Instance	
SQL Server on Azure Virtual Machines	
An Azure SQL Database single database	

Resiliency option:

	▼
Auto-failover group	
Active geo-replication	
Zone-redundant deployment	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: SQL Server on Azure VMs

SQL Server on Azure Virtual Machines can take advantage of Automated Backup, which regularly creates backups of your database to blob storage. You can also manually use this technique.

Box 2: Active geo-replication

Geo-replication for services such as Azure SQL Database and Cosmos DB will create secondary replicas of your data across multiple regions. While both services will automatically replicate data within the same region, geo-replication protects you against a regional outage by enabling you to fail over to a secondary region.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-on-azure-vm-iaas-what-i> <https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure>

**NEW QUESTION 35**

- (Exam Topic 5)

You have an Azure subscription that contains an instance of SQL Server on an Azure virtual machine named SQLVM1 and a user named User1. SQLVM1 hosts a database named DB1.

You need to ensure that User1 can create a scheduled task to perform a full backup of DB1. The solution must use the principle of least privilege.

Which built-in database role should you assign to User1?

- A. SQLAgentReaderRole
- B. db.owner
- C. SQLAgentOperatorRole
- D. SQLAgentUserRole

Answer: C

**NEW QUESTION 39**

- (Exam Topic 5)

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

Answer: D

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-qu>

**NEW QUESTION 42**

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 is 30 TB and has a 1-GB daily rate of change. You back up the database by using a Microsoft SQL Server Agent job that runs Transact-SQL commands. You perform a weekly full backup on Sunday, daily differential backups at 01:00, and transaction log backups every five minutes.

The database fails on Wednesday at 10:00.

Which three backups should you restore in sequence? To answer, move the appropriate backups from the list of backups to the answer area and arrange them in the correct order.

**Actions**

**Answer Area**

Monday, Tuesday, and then Wednesday differential backups

Wednesday, Tuesday, and then Monday log backups

full backup

Monday, Tuesday, and then Wednesday log backups

Wednesday, Tuesday, and then Monday differential backups

Wednesday log backups

Wednesday differential backup



- A. Mastered
- B. Not Mastered

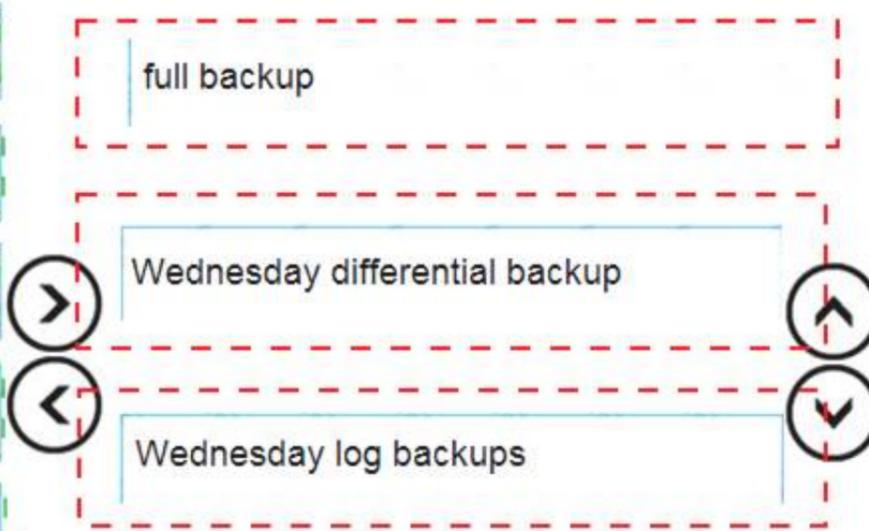
Answer: A

**Explanation:**

**Actions**

- Monday, Tuesday, and then Wednesday differential backups
- Wednesday, Tuesday, and then Monday log backups
- full backup
- Monday, Tuesday, and then Wednesday log backups
- Wednesday, Tuesday, and then Monday differential backups
- Wednesday log backups
- Wednesday differential backup

**Answer Area**



**NEW QUESTION 47**

- (Exam Topic 5)

You are building a database in an Azure Synapse Analytics serverless SQL pool. You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container. Records are structured as shown in the following sample.

```
{
  "id":123,
  "address_housenumber": "19c",
  "address_line1": "Memory Lane",
  "applicant1_name": "Jane",
  "applicant2_name": "Dev"
}
```

The records contain two applicants at most.

You need to build a table that includes only the address fields.

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.

▼ applications

CREATE EXTERNAL TABLE  
 CREATE TABLE  
 CREATE VIEW

```
WITH (
  LOCATION = 'applications/',
  DATA_SOURCE = applications_ds,
  FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
```

▼ (BULK 'https://contoso1.dfs.core.windows.net/applications/year=\*/,\*.parquet',

CROSS APPLY  
 OPENJSON  
 OPENROWSET

```
FORMAT = 'PARQUET') AS [r]
GO
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, text, application Description automatically generated

**Box 1: CREATE EXTERNAL TABLE**

An external table points to data located in Hadoop, Azure Storage blob, or Azure Data Lake Storage. External tables are used to read data from files or write data to files in Azure Storage. With Synapse SQL, you can use external tables to read external data using dedicated SQL pool or serverless SQL pool.

Syntax:

```
CREATE EXTERNAL TABLE { database_name.schema_name.table_name | schema_name.table_name | table_name }
```

```
( <column_definition> [ ,...n ] ) WITH (
```

```
LOCATION = 'folder_or_filepath', DATA_SOURCE = external_data_source_name, FILE_FORMAT = external_file_format_name
```

When using serverless SQL pool, CETAS is used to create an external table and export query results to Azure Storage Blob or Azure Data Lake Storage Gen2.

Example: AS

```
SELECT decennialTime, stateName, SUM(population) AS population FROM
```

```
OPENROWSET(BULK
```

```
'https://azureopendatastorage.blob.core.windows.net/censusdatacontainer/release/us_population_county/year=*
```

```
FORMAT='PARQUET') AS [r]
```

```
GROUP BY decennialTime, stateName GO
```

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables>

**NEW QUESTION 48**

- (Exam Topic 5)

You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container.

Which resource provider should you enable?

- A. Microsoft.EventHub
- B. Microsoft.EventGrid
- C. Microsoft.Sql
- D. Microsoft.Automation

**Answer: B**

**Explanation:**

Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.

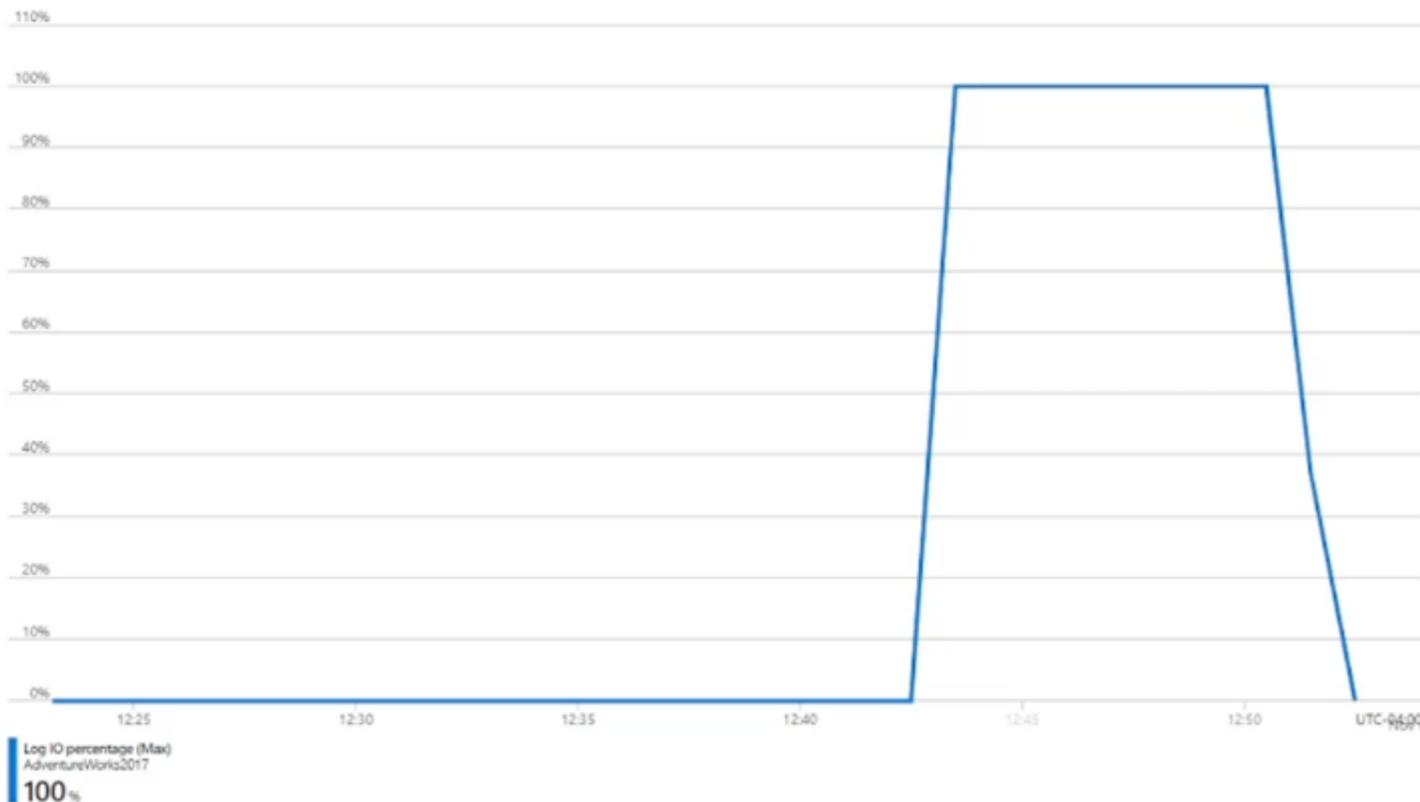
Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger>

**NEW QUESTION 52**

- (Exam Topic 5)

You have an Azure SQL database named DB1 in the General Purpose service tier. The performance metrics for DB1 are shown in the following exhibit.



You need to reduce the Log 10 percentage. The solution must minimize costs. What should you do?

- A. Increase the number of vCores.
- B. Change RecoverymodeltoSimple.
- C. Performcheckpoint operation.
- D. ChangeService tier toBusiness Critical.

**Answer: D**

**NEW QUESTION 57**

- (Exam Topic 5)

You plan to perform batch processing in Azure Databricks once daily. Which type of Databricks cluster should you use?

- A. automated
- B. interactive

C. High Concurrency

**Answer:** A

**Explanation:**

Azure Databricks makes a distinction between all-purpose clusters and job clusters. You use all-purpose clusters to analyze data collaboratively using interactive notebooks. You use job clusters to run fast and robust automated jobs.

The Azure Databricks job scheduler creates a job cluster when you run a job on a new job cluster and terminates the cluster when the job is complete.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/clusters>

**NEW QUESTION 60**

- (Exam Topic 5)

You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in five-minute intervals and report only events that arrive during the interval.

The output will be sent to a Delta Lake table. Which output mode should you use?

- A. complete
- B. append
- C. update

**Answer:** A

**Explanation:**

Complete mode: You can use Structured Streaming to replace the entire table with every batch.

Reference:

<https://docs.databricks.com/delta/delta-streaming.html>

**NEW QUESTION 64**

- (Exam Topic 5)

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 SQL Server 2019 on an Azure virtual machine.

You are troubleshooting performance issues for a query in a SQL Server instance.

To gather more information, you query `sys.dm_exec_requests` and discover that the wait type is `PAGELATCH_UP` and the `wait_resource` is `2:3:905856`.

You need to improve system performance. Solution: You create additional tempdb files. Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

**NEW QUESTION 69**

- (Exam Topic 5)

You have a data warehouse in Azure Synapse Analytics.

You need to ensure that the data in the data warehouse is encrypted at rest. What should you enable?

- A. Transparent Data Encryption (TDE)
- B. Advanced Data Security for this database
- C. Always Encrypted for all columns
- D. Secure transfer required

**Answer:** A

**Explanation:**

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. Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview>

**NEW QUESTION 73**

- (Exam Topic 5)

You have an Azure SQL database that contains a table named Customer. Customer has the columns shown in the following table.

Customer_ID	Customer_Name	Customer_Phone
11001	Contoso, Ltd.	555-555-0173
11002	Litware, Inc.	555-505-3124
11003	ADatum Corporation	555-689-4312

You plan to implement a dynamic data mask for the Customer\_Phone column. The mask must meet the following requirements:

- The first six numerals of each customer's phone number must be masked.

- > The last four digits of each customer's phone number must be visible.
- > Hyphens must be preserved and displayed.

How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

**Exposed Prefix:**  ▼  
 0  
 1  
 3  
 5

**Padding String:**  ▼  
 X  
 XXXXXX  
 XXX-XXX  
 XXX-XXX-  
 X[3]-X[3]

**Exposed Suffix:**  ▼  
 0  
 1  
 3  
 5

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: 0

Custom String : Masking method that exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix

Box 2: xxx-xxx

Box 3: 5 Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking>

**NEW QUESTION 77**

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine.

You need to add a 4-TB volume that meets the following requirements:

- > Maximizes IOPs
- > Uses premium solid state drives (SSDs)

What should you do?

- A. Attach two mirrored 4-TB SSDs.
- B. Attach a stripe set that contains four 1-TB SSDs.
- C. Attach a RAID-5 array that contains five 1-TB SSDs.
- D. Attach a single 4-TB SSD.

**Answer:** B

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/storage-configuration?tabs=window>

**NEW QUESTION 82**

- (Exam Topic 5)

You have a version-8.0 Azure Database for MySQL database.

You need to identify which database queries consume the most resources. Which tool should you use?

- A. Query Store
- B. Metrics
- C. Query Performance Insight
- D. Alerts

**Answer:** A

**Explanation:**

The Query Store feature in Azure Database for MySQL provides a way to track query performance over time. Query Store simplifies performance troubleshooting

by helping you quickly find the longest running and most resource-intensive queries. Query Store automatically captures a history of queries and runtime statistics, and it retains them for your review. It separates data by time windows so that you can see database usage patterns. Data for all users, databases, and queries is stored in the mysql schema database in the Azure Database for MySQL instance. Reference: <https://docs.microsoft.com/en-us/azure/mysql/concepts-query-store>

**NEW QUESTION 86**

- (Exam Topic 5)

You have an Azure subscription that contains a resource group named RG1. RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL. You need to use PowerShell to enable and configure automated patching for SQL. The solution must include both SQL Server and Windows security updates. How should you complete the command? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2
-MaintenanceWindowDuration 120 -PatchCategory "Important"

Get-AzVM -ResourceGroupName "RG1" -Name "SQL" |
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
-SQLManagementType
```




- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

**Answer Area**

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2
-MaintenanceWindowDuration 120 -PatchCategory "Important"

Get-AzVM -ResourceGroupName "RG1" -Name "SQL" |
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
-SQLManagementType
```




**NEW QUESTION 91**

- (Exam Topic 5)

You have an Azure SQL database.

Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue. You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort.

What should you use?

- A. Performance Recommendations in the Azure portal
- B. Extended Events in Microsoft SQL Server Management Studio (SSMS)
- C. Query Store in Microsoft SQL Server Management Studio (SSMS)
- D. Query Performance Insight in the Azure portal

**Answer:** C

**Explanation:**

Use the Query Store Page in SQL Server Management Studio.

Query performance regressions caused by execution plan changes can be non-trivial and time consuming to resolve.

Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query. This is referred to as plan forcing. Plan forcing in Query Store is provided by using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-qu>

**NEW QUESTION 95**

- (Exam Topic 5)

You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored You need to ensure that SQ Server Agent jobs running on SQL 1 can notify when a failure Occurs

Which three actions should you perform in sequence 7 TO answer. move the appropriate actions from the list Of actions to answer area and arrange them in correct order.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 98**

- (Exam Topic 5)

You have an Azure subscription that contains a SQL Server on Azure Virtual Machines instance named SQLVMI. SQLVMI hosts a database named OBI. You need to retrieve query plans from the Query Store on DBI. What should you do first?

- A. On SQLVM1, install the SQL Server IaaS Agent extension.
- B. From Microsoft SQL Server Management Studio, modify the properties of the SQL Server instance.
- C. From Microsoft SQL Server Management Studio, modify the properties of DB 1.
- D. On SQLVM1, install the Azure Monitor agent for Windows.

**Answer:** B

**NEW QUESTION 101**

- (Exam Topic 5)

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. The solution must eliminate the need for the following:

- A distributed network name (DNN)
  - A load balancer
- What should you do?

- A. Deploy VM1 and VM2 to a single proximity placement group.
- B. Deploy VM1 and VM2 to different proximity placement groups in the same Azure region.
- C. Connect VM1 and VM2 to a single subnet.
- D. Connect VM1 and VM2 to different subnets on a single virtual network.

**Answer:** D

**NEW QUESTION 103**

- (Exam Topic 5)

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 Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data

Lake Storage Gen2 container named container1.

You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1. You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.

Solution: You use an Azure Synapse Analytics serverless SQL pool to create an external table that has an additional DateTime column. Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables>

**NEW QUESTION 105**

- (Exam Topic 5)

You have a Microsoft SQL Server 2017 server that hosts five databases. You Plan to migrate the databases to Azure.

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

- > Automatically scales compute based on the workload demand
- > Provides per-second billing

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

Azure service:

Service tier:

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Azure service:

Service tier:

**NEW QUESTION 106**

- (Exam Topic 5)

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 Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that copies the data to a staging table in the data warehouse, and then uses a stored procedure to execute the R script.

Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed.

Reference:

<https://docs.microsoft.com/en-US/azure/data-factory/transform-data>

**NEW QUESTION 107**

- (Exam Topic 5)

You have an Azure SQL database named db1 on a server named server1. You need to modify the MAXDOP settings for db1. What should you do?

- A. Connect to db1 and run the sp\_configure command.
- B. Connect to the master database of server1 and run the sp\_configure command.
- C. Configure the extended properties of db1.
- D. Modify the database scoped configuration of db1.

**Answer:** D

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/configure-max-degree-of-parallelism>

**NEW QUESTION 108**

- (Exam Topic 5)

You have the following Azure Data Factory pipelines:

- > Ingest Data from System1
- > Ingest Data from System2
- > Populate Dimensions
- > Populate Facts

Ingest Data from System1 and Ingest Data from System2 have no dependencies. Populate Dimensions must execute after Ingest Data from System1 and Ingest Data from System2. Populate Facts must execute after the Populate Dimensions pipeline. All the pipelines must execute every eight hours.

What should you do to schedule the pipelines for execution?

- A. Add a schedule trigger to all four pipelines.
- B. Add an event trigger to all four pipelines.
- C. Create a parent pipeline that contains the four pipelines and use an event trigger.
- D. Create a parent pipeline that contains the four pipelines and use a schedule trigger.

**Answer:** D

**Explanation:**

Reference:

<https://www.mssqltips.com/sqlservertip/6137/azure-data-factory-control-flow-activities-overview/>

**NEW QUESTION 112**

- (Exam Topic 5)

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	Answer Area
From webapp1, connect to SQLDb1, obtain the initial dataset, and run the CHANGETABLE() function.	
Connect to SQLDb1 and run the following Transact-SQL statement.	
<code>ALTER DATABASE SQLDB1 SET CHANGE_TRACKING = ON</code>	
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.	
<code>EXEC sys.sp_cdc_enable_table</code>	
Connect to SQLDb1 and run the following Transact-SQL statement.	
<code>EXEC sys.sp_cdc_enable_db</code>	
Connect to SQLDb1 and run the following Transact-SQL statement.	
<code>ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING</code>	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, text, application Description automatically generated

**NEW QUESTION 117**

- (Exam Topic 5)

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 8 vCPUSs 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
- C. 8
- D. 64

**Answer:** C

**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>

**NEW QUESTION 122**

- (Exam Topic 5)

You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.

You need to migrate the databases hosted on SQL 1 to Azure. The solution must meet the following requirements:

The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.

Administrative effort must be minimized. What should you use to host the databases?

- A. a single Azure SQL database
- B. an Azure SQL Database elastic pool
- C. SQL Server on Azure Virtual Machines
- D. Azure SQL Managed Instance

**Answer:** D

**NEW QUESTION 125**

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool.

You run `PDW_SHOWSPACEUSED('dbo.FactInternetSales')`; and get the results shown in the following table.

ROWS	RESERVED_SPACE	DATA_SPACE	INDEX_SPACE	UNUSED_SPACE	PDW_NODE_ID	DISTRIBUTION_ID
694	2776	616	48	2112	1	1
407	2704	576	48	2080	1	2
53	2376	512	16	1848	1	3
58	2376	512	16	1848	1	4
168	2632	528	32	2072	1	5
195	2696	536	32	2128	1	6
5995	3464	1424	32	2008	1	7
0	2232	496	0	1736	1	8
264	2576	544	40	1992	1	9
3008	3016	960	32	2024	1	10
...	...	...	...	...	...	...
1550	2832	752	48	2032	1	50
1238	2832	696	40	2096	1	51
192	2632	528	32	2072	1	52
1127	2768	680	48	2040	1	53
1244	3032	704	64	2264	1	54
409	2632	568	32	2032	1	55
0	2232	496	0	1736	1	56
1437	2832	728	40	2064	1	57
0	2232	496	0	1736	1	58
384	2632	560	32	2040	1	59
225	2768	544	40	2184	1	60

Which statement accurately describes the `dbo.FactInternetSales` table?

- A. The table contains less than 10,000 rows.
- B. All distributions contain data.
- C. The table uses round-robin distribution.
- D. The table is skewed.

**Answer:** D

**Explanation:**

The rows per distribution can vary up to 10% without a noticeable impact on performance. Here the distribution varies more than 10%. It is skewed.  
 Note: `SHOWSPACEUSED` displays the number of rows, disk space reserved, and disk space used for a specific table, or for all tables in a Azure Synapse Analytics or Parallel Data Warehouse database.  
 This is a very quick and simple way to see the number of table rows that are stored in each of the 60 distributions of your database. Remember that for the most balanced performance, the rows in your distributed table should be spread evenly across all the distributions.  
`ROUND_ROBIN` distributed tables should not be skewed. Data is distributed evenly across the nodes by design.  
 Reference:  
<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribu> <https://github.com/rgl/azure-content/blob/master/articles/sql-data-warehouse/sql-data-warehouse-manage-distrib>

**NEW QUESTION 126**

- (Exam Topic 5)

You have two Azure virtual machines named VM1 and VM2 that run Windows Server 2019. VM1 and VM2 each host a default Microsoft SQL Server 2019 instance. VM1 contains a database named DB1 that is backed up to a file named `D:\DB1.bak`.

You plan to deploy an Always On availability group that will have the following configurations:

- > VM1 will host the primary replica of DB1.
- > VM2 will host a secondary replica of DB1.

You need to prepare the secondary database on VM2 for the availability group.

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

▼ DATABASE MyDB1  
BACKUP  
CREATE  
RESTORE  
 FROM DISK = 'D:\DB1.bak'  
 WITH  
 GO ▼  
NORECOVERY  
RECOVERY  
STANDBY

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, text, application, chat or text message Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondar>

**NEW QUESTION 127**

- (Exam Topic 5)

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.

**Answer:** BD

**Explanation:**

Reference:

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

**NEW QUESTION 128**

- (Exam Topic 5)

You have 20 Azure SQL databases provisioned by using the vCore purchasing model. You plan to create an Azure SQL Database elastic pool and add the 20 databases.

Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. total size of all the databases
- B. geo-replication support
- C. number of concurrently peaking databases \* peak CPU utilization per database
- D. maximum number of concurrent sessions for all the databases
- E. total number of databases \* average CPU utilization per database

**Answer:** ACE

**Explanation:**

CE: Estimate the vCores needed for the pool as follows:

For vCore-based purchasing model:  $\text{MAX}(\text{Total number of DBs} \times \text{average vCore utilization per DB},$

$\text{Number of concurrently peaking DBs} \times \text{Peak vCore utilization per DB})$

A: Estimate the storage space needed for the pool by adding the number of bytes needed for all the databases in the pool.

Reference:

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

**NEW QUESTION 129**

- (Exam Topic 5)

You have an Azure subscription.

You plan to migrate 10 on-premises Microsoft SQL Server instances to Azure.

You need to ensure that the migrated environment can be managed by using multiserver administration and supports master/target (MSX/TSX) jobs. The solution must minimize administrative effort.

Which SQL deployment options should you select as the master server (MSX) and the target server (TSX)? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

## Answer Area

MSX:

	▼
SQL database	
SQL managed instances	
SQL virtual machines	

TSX:

	▼
SQL database	
SQL managed instances	
SQL virtual machines	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

## Answer Area

MSX:

	▼
SQL database	
SQL managed instances	
SQL virtual machines	

TSX:

	▼
SQL database	
SQL managed instances	
SQL virtual machines	

### NEW QUESTION 132

- (Exam Topic 5)

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

The Intelligent Insights diagnostics log identifies queries that cause performance issues due to tempDB contention.

You need to resolve the performance issues. What should you do?

- A. Implement memory-optimized tables.
- B. Run the dbcc flushprocindbcommand.
- C. Replace the sequential index keys with nonsequential keys.
- D. Run the dbcc dbreindexcommand.

**Answer:** A

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-troubleshoot-performance#tempdb>

### NEW QUESTION 134

- (Exam Topic 5)

You have an Azure subscription that is linked to a hybrid Azure Active Directory (Azure AD) tenant. The subscription contains an Azure Synapse Analytics SQL pool named Pool1.

You need to recommend an authentication solution for Pool1. The solution must support multi-factor authentication (MFA) and database-level authentication.

Which authentication solution or solutions should you include in the recommendation? To answer, select the appropriate options in the answer area.  
 NOTE: Each correct selection is worth one point.

MFA:

▼

Azure AD authentication

Microsoft SQL Server authentication

Passwordless authentication

Windows authentication

Database-level authentication:

▼

Application roles

Contained database users

Database roles

Microsoft SQL Server logins

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Graphical user interface, text, application, chat or text message Description automatically generated

Box 1: Azure AD authentication

Azure Active Directory authentication supports Multi-Factor authentication through Active Directory Universal Authentication.

Box 2: Contained database users

Azure Active Directory Uses contained database users to authenticate identities at the database level. Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-authentication>

**NEW QUESTION 137**

- (Exam Topic 5)

You have an Azure Stream Analytics job.

You need to ensure that the job has enough streaming units provisioned. You configure monitoring of the SU % Utilization metric.

Which two additional metrics should you monitor? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Late Input Events
- B. Out of order Events
- C. Backlogged Input Events
- D. Watermark Delay
- E. Function Events

**Answer: CD**

**Explanation:**

To react to increased workloads and increase streaming units, consider setting an alert of 80% on the SU Utilization metric. Also, you can use watermark delay and backlogged events metrics to see if there is an impact.

Note: Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently non-zero, you should scale out your job, by increasing the SUs.

Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring>

**NEW QUESTION 141**

- (Exam Topic 5)

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

- Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.
- Minimize the possibility of VM1 exfiltrating data stored in SqlDb1. What should you create on VNet1?

- A. a VPN gateway
- B. a service endpoint
- C. a private link
- D. an ExpressRoute gateway

**Answer: C**

**Explanation:**

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

Reference:

<https://docs.microsoft.com/en-us/azure/private-link/private-link-overview>

#### NEW QUESTION 146

- (Exam Topic 5)

You are designing a streaming data solution that will ingest variable volumes of data. You need to ensure that you can change the partition count after creation. Which service should you use to ingest the data?

- A. Azure Event Hubs Standard
- B. Azure Stream Analytics
- C. Azure Data Factory
- D. Azure Event Hubs Dedicated

**Answer:** D

#### Explanation:

The partition count for an event hub in a dedicated Event Hubs cluster can be increased after the event hub has been created.

Reference:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions>

#### NEW QUESTION 148

- (Exam Topic 5)

You have an Azure SQL managed instance named SQLMI1 that hosts 10 databases.

You need to implement alerts by using Azure Monitor. The solution must meet the following requirements: > Minimize costs.

> Aggregate Intelligent Insights telemetry from each database. What should you do?

- A. From the Diagnostic settings of each database, select Send to Log Analytics.
- B. From the Diagnostic settings of each database, select Stream to an event hub.
- C. From the Diagnostic settings of SQLMI1, select Send to Log Analytics.
- D. From the Diagnostic settings of SQLMI1, select Stream to an event hub.

**Answer:** A

#### Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-expo>

#### NEW QUESTION 151

- (Exam Topic 5)

You have an on-premises Microsoft SQL server that uses the FileTables and FileStream features. You plan to migrate to Azure SQL.

Which service should you use?

- A. Azure SQL Database
- B. SQL Server on an Azure Virtual Machine
- C. Azure SQL Managed Instance
- D. Azure Database for MySQL

**Answer:** B

#### Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/migration-guides/database/sql-server-to-sql-database-overview>

#### NEW QUESTION 154

- (Exam Topic 5)

You have an Azure subscription that contains a logical SQL server named Server1. The master database of Server1 contains a user named User1. You need to ensure that User1 can create databases on Server1. Which database role should you assign to User1?

- A. db\_owner
- B. dbmanager
- C. dbo
- D. db\_ddladmin

**Answer:** B

#### NEW QUESTION 158

- (Exam Topic 5)

You have an Azure SQL Database elastic pool that contains 10 databases. You receive the following alert.

Msg 1132, Level 16, State 1, Line 1

The elastic pool has reached its storage limit. The storage used for the elastic pool cannot exceed (76800) MBs.

You need to resolve the alert. The solution must minimize administrative effort.

Which three actions can you perform? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Delete data from a database.
- B. Remove a database from the pool.
- C. Increase the maximum storage of the elastic pool.
- D. Shrink individual databases.
- E. Enable data compression.

**Answer:** BCD

**NEW QUESTION 162**

- (Exam Topic 5)

You configure backups for an Azure SQL database as shown in the following exhibit.

**Point-in-time-restore**

Specify how long you want to keep your point-in-time backups. [Learn more](#)

How many days would you like PITR backups to be kept? ⓘ



**Long-term retention**

Specify how long you want to keep your long-term retention backups. You may choose to keep yearly backups for up to 10 years. [Learn more](#)

**Weekly LTR Backups**

Keep weekly backups for:

**Monthly LTR Backups**

Keep the first backup of each month for:

**Yearly LTR Backups**

Keep an annual backup for:

Which weekly backup of the year would you like to keep?

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.

**Answer Area**

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a **[answer choice]**

- point-time restore (PITR) backup.
- yearly long-term retention (LTR) backup.
- weekly long-term retention (LTR) backup.
- monthly long-term retention (LTR) backup.

After the 52nd weekly backup runs, there will be **[answer choice]** in long term retention.

- 1 backup copy
- 52 backup copies
- 64 backup copies
- 65 backup copies

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice]

- point-time restore (PITR) backup.
- yearly long-term retention (LTR) backup.
- weekly long-term retention (LTR) backup.
- monthly long-term retention (LTR) backup.

After the 52nd weekly backup runs, there will be [answer choice] in long term retention.

- 1 backup copy
- 52 backup copies
- 64 backup copies
- 65 backup copies

**NEW QUESTION 166**

- (Exam Topic 5)

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)

Source spid1595

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

You need to ensure That db1 can process transactions.

**Actions**

**Answer Area**

- Add db1 back to the availability group.
- Shrink db1.
- Shrink the transaction log file.
- Remove db1 from the availability group.
- Back up the transaction log file.



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

**Actions**

**Answer Area**

- Add db1 back to the availability group.
- Shrink db1.
- Shrink the transaction log file.
- Remove db1 from the availability group.
- Back up the transaction log file.

- Remove db1 from the availability group.
- Shrink the transaction log file.
- Add db1 back to the availability group.

**NEW QUESTION 168**

- (Exam Topic 5)

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Configuration
DB1	Azure SQL Database	Hyperscale service tier No secondary replicas
App1	Azure Web Apps	App1 has read-only access to DB1. There are multiple instances of App1.

You need to create a read-only replica of DB1 and configure the App1 instances to use the replica. What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

To add read-only replicas of DB1:

▼

Create a replica on the same logical server.

Create a new logical server and configure geo-replication.

Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:

▼

Add an ApplicationIntent entry to the connection string.

Add a MultiSubnetFailover entry to the App1 connection string.

Create a dedicated endpoint and configure the App1 connection string to point to the endpoint.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Text Description automatically generated

Reference:

<https://sqlserverguides.com/read-only-replica-azure-sql/>

**NEW QUESTION 172**

- (Exam Topic 5)

You have an Azure SQL database named DB1. DB1 has a table named Table1 that contains the following columns.

Name	Type
Column1	Ntext
Column2	Geometry
Column3	Image
Column4	Varchar
Column5	Datetime2

You plan to enable Always Encrypted for Table1.

Which two columns support encryption? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point

- A. Column1
- B. Column2
- C. Column3
- D. Column4
- E. Column5

**Answer:** AD

**NEW QUESTION 173**

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and a database named DB1. DB1 contains a fact table named Table.

You need to identify the extent of the data skew in Table1. What should you do in Synapse Studio?

- A. Connect to Pool1 and query sys.dm\_pdw\_nodes\_db\_partition\_stats.
- B. Connect to the built-in pool and run DBCC CHECKALLOC.
- C. Connect to Pool1 and run DBCC CHECKALLOC.
- D. Connect to the built-in pool and query sys.dm\_pdw\_nodes\_db\_partition\_stats.

**Answer:** D

**Explanation:**

Use sys.dm\_pdw\_nodes\_db\_partition\_stats to analyze any skewness in the data. Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/cheat-sheet>

**NEW QUESTION 177**

- (Exam Topic 5)

You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL. Which switch should you use to switch between languages?

- A. \[<language>]
- B. %<language>
- C. \[<language>]
- D. @<language>

**Answer: B**

**Explanation:**

You can override the default language by specifying the language magic command %<language> at the beginning of a cell. The supported magic commands are: %python, %r, %scala, and %sql.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/notebooks/notebooks-use>

**NEW QUESTION 178**

- (Exam Topic 5)

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.

You plan to migrate DB1 to an Azure SQL Database managed instance.

What should you use to minimize downtime and data loss during the migration?

- A. distributed availability groups
- B. database mirroring
- C. log shipping
- D. Database Migration Assistant

**Answer: D**

**Explanation:**

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

**NEW QUESTION 179**

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2016 instance that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named Mil.

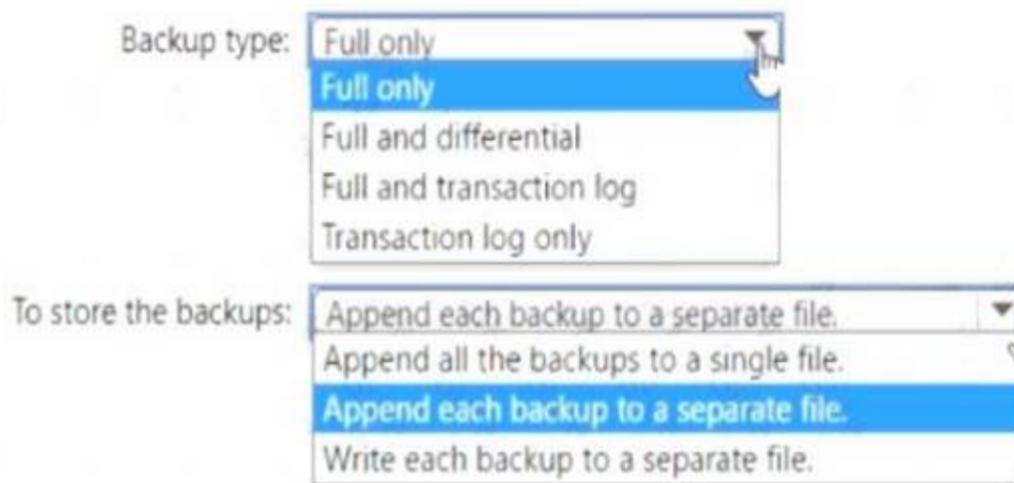
You plan to perform an online migration of db1 to MM by using Azure Database Migration Service.

You need to create the backups for the migration. The solution must minimize the number of backup files created.

Which type of backups should you create, and how should you store the backups? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

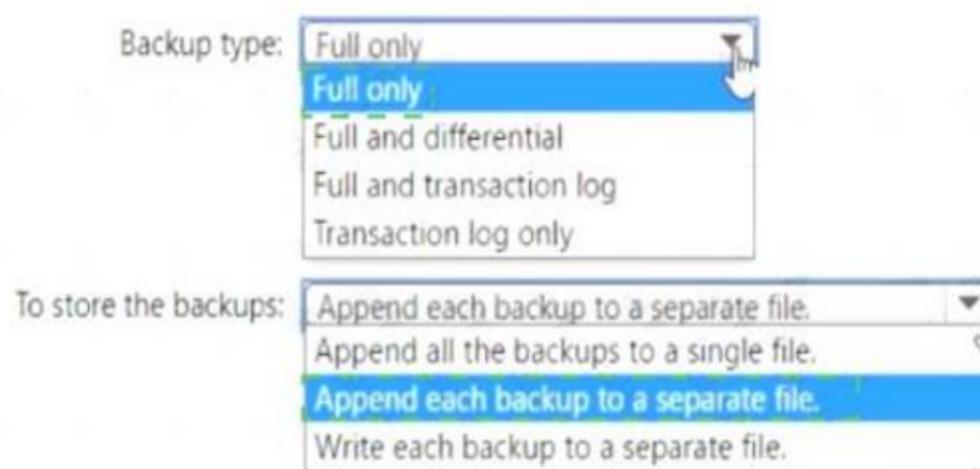


- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Answer Area



**NEW QUESTION 182**

- (Exam Topic 5)

You have an Azure Synapse Analytics workspace named WS1 that contains an Apache Spark pool named Pool1.

You plan to create a database named DB1 in Pool1.

You need to ensure that when tables are created in DB1, the tables are available automatically as external tables to the built-in serverless SQL pool.

Which format should you use for the tables in DB1?

- A. JSON
- B. CSV
- C. Parquet
- D. ORC

**Answer: C**

**Explanation:**

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.

For each Spark external table based on Parquet and located in Azure Storage, an external table is created in a serverless SQL pool database. As such, you can shut down your Spark pools and still query Spark external tables from serverless SQL pool.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-storage-files-spark-tables>

**NEW QUESTION 184**

- (Exam Topic 5) You have an Azure SQL database. You identify a long running query.

You need to identify which operation in the query is causing the performance issue.

What should you use to display the query execution plan in Microsoft SQL Server Management Studio (SSMS)?

- A. Live Query Statistics
- B. an estimated execution plan
- C. an actual execution plan
- D. Client Statistics

**Answer: A**

**Explanation:**

<https://www.mssqltips.com/sqlservertip/3685/live-query-statistics-in-sql-server-2016/>

**NEW QUESTION 189**

- (Exam Topic 5)

You are designing a security model for an Azure Synapse Analytics dedicated SQL pool that will support multiple companies.

You need to ensure that users from each company can view only the data of their respective company. Which two objects should you include in the solution? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. a column encryption key
- B. asymmetric keys
- C. a function
- D. a custom role-based access control (RBAC) role
- E. a security policy

**Answer: DE**

**Explanation:**

Azure RBAC is used to manage who can create, update, or delete the Synapse workspace and its SQL pools, Apache Spark pools, and Integration runtimes. Define and implement network security configurations for resources related to your dedicated SQL pool with Azure Policy.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/security/synapse-workspace-synapse-rbac> <https://docs.microsoft.com/en-us/security/benchmark/azure/baselines/synapse-analytics-security-baseline>

**NEW QUESTION 190**

- (Exam Topic 5)

You have an Azure SQL database that contains a table named Employees. Employees contains a column named Salary. You need to encrypt the Salary column. The solution must prevent database administrators from reading the data in the Salary column and must provide the most secure encryption. 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	Answer Area
Encrypt the Salary column by using the randomized encryption type.	
Create a column encryption key.	
Enable Transparent Data Encryption (TDE).	 
Encrypt the Salary column by using the deterministic encryption type.	 
Apply a dynamic data mask to the Salary column.	
Create a column master key.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create a column master key

Create a column master key metadata entry before you create a column encryption key metadata entry in the database and before any column in the database can be encrypted using Always Encrypted.

Step 2: Create a column encryption key.

Step 3: Encrypt the Salary column by using the randomized encryption type.

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.

Note: A column encryption key metadata object contains one or two encrypted values of a column encryption key that is used to encrypt data in a column. Each value is encrypted using a column master key.

Reference:

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

**NEW QUESTION 195**

- (Exam Topic 5)

You have an Azure subscription.

You plan to deploy an Azure SQL database by using an Azure Resource Manager template.

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

NOTE: Each correct selection is worth one point.

**Answer Area**

```
{
  "resources": [
    {
      "type": 
      "apiVersion": "2020-02-02-preview",
      "name": "[parameters('name1')]",
      "location": "[parameters('location')]",
      ...
      "resources": [
        {
          "type": "databases",
          "apiVersion": "2020-02-02-preview",
          ...
          "dependsOn": [
            "properties": [
              "tags": [
                "[resourceId('Microsoft.Sql/servers', concat(parameters('name1')))]"
              ]
            ]
          ]
        }
      ]
    }
  ]
}
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Text Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart>

**NEW QUESTION 199**

- (Exam Topic 5)

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
SQL1	SQL Server on Azure Virtual Machines	Not applicable
db1	Microsoft SQL Server database	Hosted on SQL1
mysqlbackups	General purpose v2 storage account	Not applicable

You need to back up db1 to mysqlbackups, and then restore the backup to a new database named db2 that is hosted on SQL1. The solution must ensure that db1 is backed up to a stripe set.

Which three Transact-SQL statements should you execute in sequence? To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.

**Statements**

**Answer Area**

```
RESTORE DATABASE db2 FROM URL = URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup', RECOVERY,
MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup';
GO
```

```
RESTORE DATABASE db2 FROM URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_4.bak'
WITH RECOVERY,
MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

```
CREATE CREDENTIAL
[https://mysqlbackups.blob.core.windows.net/backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = '<SAS_TOKEN>'
GO
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_4.bak'
GO
```

```
CREATE CREDENTIAL [sqlbackup] WITH IDENTITY
=
'sqlsamplebackup'
, SECRET = '<mystorageaccountaccesskey>';
GO
```



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-to-url?view=sql-serv>

**NEW QUESTION 202**

- (Exam Topic 5)

You have an Azure SQL database named DB1.

You need to display the estimated execution plan of a query by using the query editor in the Azure portal.

What should you do first?

- A. Run the `showplan_all` Transact-SQL statement.
- B. For DB1, set `QUERY_CAPTURE_MODE` of Query Store to All.
- C. Run the `forceplan` Transact-SQL statement.
- D. Enable Query Store for DB1.

**Answer:** A

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-showplan-all-transact-sql?view=sql-server-ver15>

**NEW QUESTION 207**

- (Exam Topic 5)

You have an instance of SQL Server on Azure Virtual Machines named SQL1.

SQL1 contains an Extended Events session named session1 that captures Microsoft SQL Server events. You need to correlate the session events with events captured by Event Tracing for Windows (ETW). What should you do for session1?

- A. Modify the Set Session Event Filters settings.
- B. Add a target.
- C. Add an action.
- D. Modify the Specify Session Data Storage settings.

**Answer:** B

**NEW QUESTION 209**

- (Exam Topic 5)

You are developing an application that uses Azure Data Lake Storage Gen 2.

You need to recommend a solution to grant permissions to a specific application for a limited time period. What should you include in the recommendation?

- A. role assignments
- B. account keys
- C. shared access signatures (SAS)
- D. Azure Active Directory (Azure AD) identities

**Answer:** C

**Explanation:**

A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:

What resources the client may access.

What permissions they have to those resources. How long the SAS is valid.

Note: Data Lake Storage Gen2 supports the following authorization mechanisms:

- > Shared Key authorization
- > Shared access signature (SAS) authorization
- > Role-based access control (Azure RBAC)
- > Shared Key authorization
- > Shared access signature (SAS) authorization
- > Role-based access control (Azure RBAC)
- > Access control lists (ACL)

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

**NEW QUESTION 214**

- (Exam Topic 5)

You need to apply 20 built-in Azure Policy definitions to all new and existing Azure SQL Database deployments in an Azure subscription. The solution must minimize administrative effort.

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	Answer Area
Duplicate Azure Policy definitions	
Run Azure Policy remediation tasks	
Create an Azure Blueprints assignment	⬆
Create an Azure Policy initiative	⬇
Create an Azure Policy initiative assignment	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create an Azure Policy Initiative

The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take.

With an initiative definition, you can group several policy definitions to achieve one overarching goal. An initiative evaluates resources within scope of the assignment for compliance to the included policies.

Step 2: Create an Azure Policy Initiative assignment

Assign the initiative definition you created in the previous step. Step 3: Run Azure Policy remediation tasks

To apply the Policy Initiative to the existing SQL databases. Reference:

<https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage>

**NEW QUESTION 218**

- (Exam Topic 5)

You have an Azure SQL database named sqldb1.

You need to minimize the amount of space by the data and log files of sqldb1. What should you run?

- A. DBCC SHRINKDATABASE
- B. sp\_clean\_db\_free\_space
- C. sp\_clean\_db\_file\_free\_space
- D. DBCC SHRINKFILE

**Answer: D**

**Explanation:**

DBCC SHRINKDATABASE shrinks the size of the data and log files in the specified database. Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkdatabase-transact-sql>

**NEW QUESTION 223**

- (Exam Topic 5)

You have an Azure SQL logical server. You run the following script.

```
CREATE DATABASE Sales
GO
CREATE TABLE [dbo].[Orders]
(
    [OrderID] INT NOT NULL,
    [OrderDescription] NVARCHAR (MAX) NOT NULL,
    [Timestamp] Datetime2 NOT NULL
)
WITH (
    SYSTEM_VERSIONING = ON,
    LEDGER = ON
);
GO
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

**NOTE:** Each correct selection is worth one point

Statements	Yes	No
The <code>orders</code> table will allow only rows to be inserted.	<input type="radio"/>	<input type="radio"/>
To create additional tables in the <code>Sales</code> database, the <code>LEDGER = ON</code> parameter must be used.	<input type="radio"/>	<input type="radio"/>
To ensure that a timestamp is added to each row in the <code>orders</code> table, the <code>GENERATED ALWAYS</code>	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Yes

No No

**NEW QUESTION 227**

- (Exam Topic 5)

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 SQL Server 2019 on an Azure virtual machine.

You are troubleshooting performance issues for a query in a SQL Server instance.

To gather more information, you query sys.dm\_exec\_requests and discover that the wait type is PAGELATCH\_UP and the wait\_resource is 2:3:905856.

You need to improve system performance. Solution: You shrink the transaction log file. Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

**NEW QUESTION 231**

- (Exam Topic 5)

You have an Azure SQL database named DB1.

You need to encrypt DB1. The solution must meet the following requirements;

- Encrypt data in motion.
- Support comparison operators.
- Provide randomized encryption.

What should you include in the solution?

- A. Always Encrypted
- B. column-level encryption
- C. Transparent Data Encryption (TDE)
- D. Always Encrypted with secure enclaves

**Answer: A**

**NEW QUESTION 235**

- (Exam Topic 5)

You have an Azure subscription.

You plan to deploy an instance of SQL Server on Azure Virtual Machines that supports Write Accelerator.

Which virtual machine series should you use?

- A. H-series
- B. G-series
- C. M-series
- D. E-series

**Answer: C**

**NEW QUESTION 240**

- (Exam Topic 5)

You have an Azure Data Factory pipeline that is triggered hourly. The pipeline has had 100% success for the past seven days.

The pipeline execution fails, and two retries that occur 15 minutes apart also fail. The third failure returns the following error.

```
ErrorCode=UserErrorFileNotFound,  
'Type=Microsoft.DataTransfer.Common.Shared.HybridDeliveryException,Message=ADLS  
Gen2 operation failed for: Operation returned an invalid status code  
'NotFound'. Account: 'contosoproduksouth' FileSystem: wwi.Path:  
'BIKES/CARBON/year=2021/month=01/day=10/hour=06'. ErrorCode:  
'PathNotFound'.Message: 'The specified path does not exist.'.RequestId:  
'6d269b78-901f-001b-4924-e7a7bc000000'. TimeStamp: 'Sun, 10 Jan 2021 07:45:05
```

What is a possible cause of the error?

- A. From 06:00 to 07:00 on January 10, 2021, there was no data in wwi/BIKES/CARBON.
- B. The parameter used to generate year=2021/month=01/day=10/hour=06 was incorrect.
- C. From 06:00 to 07:00 on January 10, 2021, the file format of data in wwi/BIKES/CARBON was incorrect.
- D. The pipeline was triggered too early.

**Answer: B**

**NEW QUESTION 241**

- (Exam Topic 5)

You have an Azure SQL database named db1 on a server named server1. You use Query Performance Insight to monitor db1.

You need to modify the Query Store configuration to ensure that performance monitoring data is available as soon as possible.

Which configuration setting should you modify and which value should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS
INTERVAL_LENGTH_MINUTES
MAX_PLANS_PER_QUERY
QUERY_CAPTURE_MODE

Value:

1
60
CUSTOM
ON

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, text, application Description automatically generated

**NEW QUESTION 245**

- (Exam Topic 5)

You need to use an Azure Resource Manager (ARM) template to deploy an Azure virtual machine that will host a Microsoft SQL Server instance. The solution must maximize disk I/O performance for the SQL Server database and log files

How should you complete the template? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```

"variables": {
  "dataDisks": {
    "caching":  "dataDiskCount": 8, "logDisksCount": 1,
    ...
  }
}

"resources": [
  ...
  "osDisk": {
    ...
    "copy": [
      {
        "name": "dataDisks", "count": "[add(variables('dataDiskCount'), variables('logDisksCount'))]",
        "input": {
          "lun": "[copyIndex('dataDisks')]", "createOption": "empty",
          "caching": "[if(greaterOrEquals(copyIndex('dataDisks'), parameters('dataDiskCount')), , variables('dataDisks').caching )]", "diskSizeGB": 1023,
        }
      }
    ]
  }
]

```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Read only ReadWrite

**NEW QUESTION 249**

- (Exam Topic 5)

You have several Azure SQL databases on the same Azure SQL Database server in a resource group named ResourceGroup1.

You must be alerted when CPU usage exceeds 80 percent for any database. The solution must apply to any additional databases that are created on the Azure SQL server.

Which resource type should you use to create the alert?

- A. Resource Groups
- B. SQL Servers
- C. SQL Databases

D. SQL Virtual Machines

**Answer:** C

**Explanation:**

There are resource types related to application code, compute infrastructure, networking, storage + databases. You can deploy up to 800 instances of a resource type in each resource group.

Some resources can exist outside of a resource group. These resources are deployed to the subscription, management group, or tenant. Only specific resource types are supported at these scopes.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/resource-providers-and-types>

**NEW QUESTION 254**

- (Exam Topic 5)

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.

You plan to migrate DB1 to an Azure SQL Database managed instance.

What should you use to minimize downtime and data loss during the migration?

- A. database mirroring
- B. distributed availability groups
- C. Always On Availability Group
- D. Azure Database Migration Service

**Answer:** D

**NEW QUESTION 255**

- (Exam Topic 5)

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

**Answer:** D

**Explanation:**

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>

**NEW QUESTION 259**

- (Exam Topic 5)

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.

**Answer:** D

**Explanation:**

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

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

**NEW QUESTION 264**

- (Exam Topic 5)

You have an Azure SQL database named sqldb1.

You need to minimize the possibility of Query Store transitioning to a read-only state. What should you do?

- A. Halve the value of Data Flush Interval.
- B. Double the value of Statistics Collection Interval.
- C. Halve the value of Statistics Collection Interval.
- D. Double the value of Data Flush Interval.

**Answer:** A

**Explanation:**

The Max Size (MB) limit isn't strictly enforced. Storage size is checked only when Query Store writes data to disk. This interval is set by the Data Flush Interval (Minutes) option. If Query Store has breached the

Maximum size limit between storage size checks, it transitions to read-only mode. Reference:  
<https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store>

**NEW QUESTION 265**

- (Exam Topic 5)

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_ERRORMSG
NO_INFOMSGS
PHYSICAL_ONLY

)

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Table Description automatically generated with low confidence

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql?view=sql-ser>

**NEW QUESTION 269**

- (Exam Topic 5)

You need to recommend an availability strategy for an Azure SQL database. The strategy must meet the following requirements:

- Support failovers that do not require client applications to change their connection strings.
- Replicate the database to a secondary Azure region.
- Support failover to the secondary region. What should you include in the recommendation?

- A. failover groups
- B. transactional replication
- C. Availability Zones
- D. geo-replication

**Answer:** A

**Explanation:**

Active geo-replication is an Azure SQL Database feature that allows you to create readable secondary databases of individual databases on a server in the same or different data center (region).

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview>

**NEW QUESTION 271**

- (Exam Topic 5)

You have an Azure subscription.

You need to deploy an Azure SQL managed instance that meets the following requirements:

- Optimize latency.
- Maximize the memory-to-vCore ratio.

Which service tier and hardware generation should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Service tier: 

▼
Business Critical
General Purpose
Hyperscale

Hardware generation: 

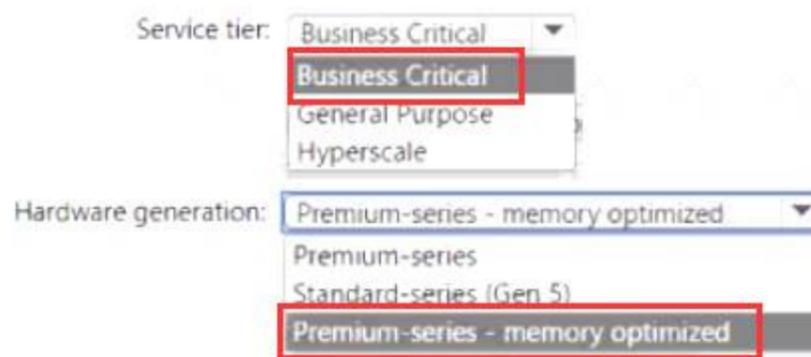
▼
Premium-series - memory optimized
Premium-series
Standard-series (Gen 5)
Premium-series - memory optimized

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area



**NEW QUESTION 276**

- (Exam Topic 5)

You have an Azure Data Factory pipeline that performs an incremental load of source data to an Azure Data Lake Storage Gen2 account. Data to be loaded is identified by a column named LastUpdatedDate in the source table. You plan to execute the pipeline every four hours. You need to ensure that the pipeline execution meets the following requirements:  
 Automatically retries the execution when the pipeline run fails due to concurrency or throttling limits. Supports backfilling existing data in the table.  
 Which type of trigger should you use?

- A. tumbling window
- B. on-demand
- C. event
- D. schedule

**Answer:** A

**Explanation:**

The Tumbling window trigger supports backfill scenarios. Pipeline runs can be scheduled for windows in the past.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipeline-execution-triggers>

**NEW QUESTION 281**

- (Exam Topic 4)

You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements. Which Azure Storage functionality should you include in the solution?

- A. time-based retention
- B. change feed
- C. lifecycle management
- D. soft delete

**Answer:** C

**Explanation:**

The lifecycle management policy lets you:

Delete blobs, blob versions, and blob snapshots at the end of their lifecycles Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts>

**NEW QUESTION 284**

- (Exam Topic 4)

You need to implement the surrogate key for the retail store table. The solution must meet the sales transaction dataset requirements. What should you create?

- A. a table that has a FOREIGN KEY constraint
- B. a table the has an IDENTITY property
- C. a user-defined SEQUENCE object
- D. a system-versioned temporal table

**Answer:** B

**Explanation:**

Scenario: Contoso requirements for the sales transaction dataset include: Implement a surrogate key to account for changes to the retail store addresses.

A surrogate key on a table is a column with a unique identifier for each row. The key is not generated from the table data. Data modelers like to create surrogate keys on their tables when they design data warehouse models. You can use the IDENTITY property to achieve this goal simply and effectively without affecting load performance.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tablesidentity>

**NEW QUESTION 285**

- (Exam Topic 2)

Based on the PaaS prototype, which Azure SQL Database compute tier should you use?

- A. Business Critical 4-vCore
- B. Hyperscale
- C. General Purpose v-vCore

D. Serverless

**Answer:** A

**Explanation:**

There are CPU and Data I/O spikes for the PaaS prototype. Business Critical 4-vCore is needed. Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview>

**NEW QUESTION 288**

- (Exam Topic 1)

You need to recommend the appropriate purchasing model and deployment option for the 30 new databases. The solution must meet the technical requirements and the business requirements.

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

NOTE:Each correct selection is worth one point.

Purchasing model:  ▼

- Azure virtual machine reserved instances
- DTU
- vCore

Deployment option:  ▼

- An Azure SQL Database elastic pool
- An Azure SQL Database managed instance
- A SQL Server Always On availability group

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: DTU

Scenario:

- > The 30 new databases must scale automatically.
- > Once all requirements are met, minimize costs whenever possible.

You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.

In short, for simplicity, the DTU model has an advantage. Plus, if you're just getting started with Azure SQL Database, the DTU model offers more options at the lower end of performance, so you can get started at a lower price point than with vCore.

Box 2: An Azure SQL database elastic pool

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. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

Reference:

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

**NEW QUESTION 293**

- (Exam Topic 1)

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

- 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:** B

**Explanation:**

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

**NEW QUESTION 297**

- (Exam Topic 1)

You create all of the tables and views for ResearchDB1.

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

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

- Run the Always Encrypted wizard.
- Create an Azure Key Vault instance and generate a secret.
- Create an Azure Key Vault instance and configure an access policy.
- Create an Azure AD managed identity.
- Register ResearchApp1 to Azure AD.

### Answer Area



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, text, application Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/always-encrypted-azure-key-vault-configure?tabs=az>

**NEW QUESTION 298**

- (Exam Topic 1)

You need to recommend a solution to ensure that the customers can create the database objects. The solution must meet the business goals. What should you include in the recommendation?

- A. For each customer, grant the customer ddl\_admin to the existing schema.
- B. For each customer, create an additional schema and grant the customer ddl\_admin to the new schema.
- C. For each customer, create an additional schema and grant the customer db\_writerto the new schema.
- D. For each customer, grant the customer db\_writerto the existing schema.

**Answer:** B

**NEW QUESTION 302**

- (Exam Topic 1)

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

**Answer:** AE

**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 Reference:

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

**NEW QUESTION 304**

- (Exam Topic 1)

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	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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 publicVirtual 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/)

**NEW QUESTION 308**

.....

## Relate Links

**100% Pass Your DP-300 Exam with ExamBible Prep Materials**

<https://www.exambible.com/DP-300-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>