Microsoft

Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)

NEW QUESTION 1
- (Exam Topic 5)
You have an Azure SQL Database managed instance named sqldbmi1 that contains a database name Sales. You need to initiate a backup of Sales.
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.

BACKUP DATABASE Sales

▼
TO DISK = \\BackupSystem\BackupDisk1\Sales.bak'
TO DISK = 'X:\BAK\Sales.bak'
TO 'Sales_Backup'
TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak'

WITH STATS = 5,

▼
WITH COPY_ONLY;
WITH ENCRYPTION;
WITH FILE_SNAPSHOT;
WITH NO_TRUNCATE

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak' Native database backup in Azure SQL Managed Instance.
You can backup any database using standard BACKUP T-SQL command: BACKUP DATABASE tpcc2501
TO URL = 'https://myacc.blob.core.windows.net/testcontainer/tpcc2501.bak'
WITH COPY_ONLY
Box 2: WITH COPY_ONLY
Reference:
https://techcommunity.microsoft.com/t5/azure-sql-database/native-database-backup-in-azure-sql-managed-insta

NEW QUESTION 2
- (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 Business Critical service tier and Availability Zones.
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
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 3
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1. The database reports a CHECKSUM error.
You need to recover the database.
How should you complete the statements? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
USE master;

ALTER DATABASE [DB1] SET  [          ▼]  WITH ROLLBACK IMMEDIATE;
                          OFFLINE
GO                        ONLINE
                          SINGLE_USER
                          TRUSTWORTHY


DBCC CHECKDB ('DB1',      [          ▼]  WITH NO_INFOMSGS;
                          MOINDEX
GO                        PHYSICAL_ONLY
                          REPAIR_ALLOW_DATA_LOSS
                          REPAIR_FAST


ALTER DATABASE [DB1] SET  [          ▼]
                          MULTI_USER;
GO                        ONLINE;
                          OPEN;
                          TRUSTWORTHY;
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: SINGLE_USER
The specified database must be in single-user mode to use one of the following repair options. Box 2: REPAIR_ALLOW_DATA_LOSS
REPAIR_ALLOW_DATA_LOSS tries to repair all reported errors. These repairs can cause some data loss.
Note: The REPAIR_ALLOW_DATA_LOSS option is a supported feature but it may not always be the best option for bringing a database to a physically consistent state. If successful, the REPAIR_ALLOW_DATA_LOSS option may result in some data loss. In fact, it may result in more data lost than if a user were to restore the database from the last known good backup.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql

NEW QUESTION 4
- (Exam Topic 5)
You have an Azure SQL database named DB 1 in the General Purpose service tier. You need to monitor DB 1 by using SQL Insights.
What should you include in the solution? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```
To collect monitoring data, use:   [                    ▼]
                                   A virtual machine
                                   An Azure function
                                   The Azure Monitor agent


To store monitoring data, create:  [                    ▼]
                                    A Log Analytics workspace
                                    An Azure SQL database
                                    An Azure Storage account
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1 = Azure Monitor Agent Box 2 = An Azure SQL database
https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql

NEW QUESTION 5
- (Exam Topic 5)
You are creating a managed data warehouse solution on Microsoft Azure.
You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails.
You need to configure Azure Synapse Analytics to receive the data.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

Actions
Create an external data source for Azure Blob storage.
Create a master key on database.
Enable Transparent Data Encryption.
Create the external table FactSalesOrderDetails.
Load the data to a staging table.
Create an external file format to map the parquet files.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, chat or text message Description automatically generated
To query the data in your Hadoop data source, you must define an external table to use in Transact-SQL queries. The following steps describe how to configure the external table.
Step 1: Create a master key on database.
* 1. Create a master key on the database. The master key is required to encrypt the credential secret. (Create a database scoped credential for Azure blob storage.)
Step 2: Create an external data source for Azure Blob storage.
* 2. Create an external data source with CREATE EXTERNAL DATA SOURCE.. Step 3: Create an external file format to map the parquet files.
* 3. Create an external file format with CREATE EXTERNAL FILE FORMAT. Step 4. Create an external table FactSalesOrderDetails
* 4. Create an external table pointing to data stored in Azure storage with CREATE EXTERNAL TABLE. Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure-blob-storage

NEW QUESTION 6
- (Exam Topic 5)
You are building an Azure virtual machine.
You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS. You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.
You need to configure storage for the virtual machine to support the database.
Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

Actions

Answer Area

Actions
a virtual disk that uses the stripe layout
a virtual disk that uses the mirror layout
a volume
a virtual disk that uses the simple layout
a storage pool

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Follow these same steps to create striped virtual disk:

≫ Create Log Storage Pool.
≫ Create Virtual Disk
≫ Create Volume
Box 1: a storage pool
Box 2: a virtual disk that uses stripe layout
Disk Striping: Use multiple disks and stripe them together to get a combined higher IOPS and Throughput limit. The combined limit per VM should be higher than the combined limits of attached premium disks.
Box 3: a volume Reference:
https://hanu.com/hanu-how-to-striping-of-disks-for-azure-sql-server/

NEW QUESTION 7
- (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 8
- (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 9
- (Exam Topic 5)
You deploy a database to an Azure SQL Database managed instance.
You need to prevent read queries from blocking queries that are trying to write to the database. Which database option should set?

A. PARAMETERIZATION to FORCED
B. PARAMETERIZATION to SIMPLE
C. Delayed Durability to Forced
D. READ_COMMITTED_SNAPSHOT to ON

Answer: D

Explanation:
In SQL Server, you can also minimize locking contention while protecting transactions from dirty reads of uncommitted data modifications using either:

› The READ COMMITTED isolation level with the READ_COMMITTED_SNAPSHOT database option set to ON.

› The SNAPSHOT isolation level.

If READ_COMMITTED_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/set-transaction-isolation-level-transact-sql

NEW QUESTION 10
- (Exam Topic 5)
You have an on-premises multi-tier application named App1 that includes a web tier, an application tier, and a Microsoft SQL Server tier. All the tiers run on Hyper-V virtual machines.
Your new disaster recovery plan requires that all business-critical applications can be recovered to Azure. You need to recommend a solution to fail over the database tier of App1 to Azure. The solution must provide the ability to test failover to Azure without affecting the current environment.
What should you include in the recommendation?

A. Azure Backup
B. Azure Information Protection
C. Windows Server Failover Cluster
D. Azure Site Recovery

Answer: D

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-test-failover-to-azure

NEW QUESTION 10
- (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":  ┌─────────────────────────────────────┐ ▼
              │ "Microsoft.Sql/servers"              │
              │ "Microsoft.SqlVirtualMachines/sqlVirtualMachines" │
              │ "Microsoft.Synapse/workspaces/sqldatabases" │
              └─────────────────────────────────────┘
      "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 14
- (Exam Topic 5)
A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

```
CREATE TABLE dbo.Employee
(
    EmployeeID      INT IDENTITY(1,1) PRIMARY KEY CLUSTERED NOT NULL,
    FirstName       VARCHAR(100) NOT NULL,
    LastName        VARCHAR(100) NOT NULL,
    Title           VARCHAR(100) NULL,
    LastHireDate    DATETIME NULL,
    StreetAddress1  VARCHAR(500) NOT NULL,
    StreetAddress2  VARCHAR(500) NOT NULL,
    StreetAddress3  VARCHAR(500) NOT NULL,
    City            VARCHAR(200) NOT NULL,
    StateName       VARCHAR(20) NOT NULL,
    Salary          VARCHAR(20) NULL,
    PhoneNumber     VARCHAR(20) NOT NULL
)
```

You need to recommend changes to the data types to reduce storage and improve performance. Which two actions should you recommend? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Change Salary to the money data type.
B. Change PhoneNumber to the float data type.
C. Change LastHireDate to the datetime2(7) data type.
D. Change PhoneNumber to the bigint data type.
E. Change LastHireDate to the date data type.

Answer: AE

NEW QUESTION 17

- (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 18
- (Exam Topic 5)
You have an Azure SQL database that contains a table named factSales. FactSales contains the columns shown in the following table.

Name	Data type
SalesID	Int
Product	Int
Total Number	Numeric(8,4)
Tax Number	Numeric(8,4)
SalesRep	Varchar(30)

FactSales has 6 billion rows and is loaded nightly by using a batch process.
Which type of compression provides the greatest space reduction for the database?

A. page compression
B. row compression
C. columnstore compression
D. columnstore archival compression

Answer: D

Explanation:
Columnstore tables and indexes are always stored with columnstore compression. You can further reduce the size of columnstore data by configuring an additional compression called archival compression.
Note: Columnstore — The columnstore index is also logically organized as a table with rows and columns, but the data is physically stored in a column-wise data format.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression

NEW QUESTION 19
- (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 21
- (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 26
- (Exam Topic 5)
You are provisioning an Azure SQL database in the Azure portal 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.

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

no extra time
up to 10 minutes
up to one minute

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

intermittent and unpredictable
regular and high
steady and low

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview

NEW QUESTION 28
- (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 29
- (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 32
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL managed instance named SQLMi1 and a SQL Agent job named Backupdb. Backupdb performs a daily backup of the databases hosted on SQLMi1.
You need to be notified by email if the job fails.
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.
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions

| Create a SQL Server Agent alert. |
| Create an operator. |
| Create an extended event. |
| Enable Database Mail. |
| Add a failure notification to the job. |

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/job-automation-managed-instance

NEW QUESTION 37
- (Exam Topic 5)
You have a new Azure subscription.
You create an Azure SQL Database instance named DB1 on an Azure SQL Database server named Server1. You need to ensure that users can connect to DB1 in the event of an Azure regional outage. In the event of an outage, applications that connect to DB1 must be able to connect without having to update the connection
strings.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. From the properties of DB1. configure geo-replication.
B. From the properties of Server1 add a failover group.
C. Create a new Azure SQL Database server named Server2.
D. From the properties of Server1 configure retention for DB1
E. Create a new Azure SQL Database instance named DB2.

Answer: BC

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview?tabs=azure-powershell https://docs.microsoft.com/en-us/azure/azure-sql/database/failover-group-add-single-database-tutorial?tabs=azur

NEW QUESTION 42
- (Exam Topic 5)
You have an Azure Active Directory (Azure AD) tenant named contoso.com that contains a user named user1@contoso.com and an Azure SQL managed instance named SQLMI1.
You need to ensure that user1@contoso.com can create logins in SQLMI1 that map to Azure AD service principals.
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
Run `CREATE LOGIN user1@contoso.com FROM EXTERNAL PROVIDER` on the master database.	
Run `ALTER SERVER ROLE securityadmin ADD MEMBER user1@contoso.com`.	⓵ ⓷
Create a managed identity for SQLMI1.	⓸
Grant SQLMI1 read access to Azure AD.	
Run `CREATE USER user1@contoso.com FROM LOGIN user1@contoso.com`.	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated with medium confidence
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/aad-security-configure-tutorial

NEW QUESTION 47
- (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 49
- (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 52
- (Exam Topic 5)
You have an Azure Data Factory instance named ADF1 and two Azure Synapse Analytics workspaces named WS1 and WS2.
ADF1 contains the following pipelines:

≫ P1:Uses a copy activity to copy data from a nonpartitioned table in a dedicated SQL pool of WS1 to an Azure Data Lake Storage Gen2 account

≫ P2:Uses a copy activity to copy data from text-delimited files in an Azure Data Lake Storage Gen2 account to a nonpartitioned table in a dedicated SQL pool of WS2
You need to configure P1 and P2 to maximize parallelism and performance.
Which dataset settings should you configure for the copy activity of each pipeline? To answer, select the appropriate options in the answer area.

P1:
| Set the Copy method to Bulk insert. |
| Set the Copy method to PolyBase. |
| Set the Isolation level to Repeatable read. |
| Set the Partition option to Dynamic range. |

P2:
| Set the Copy method to Bulk insert. |
| Set the Copy method to PolyBase. |
| Set the Isolation level to Repeatable read. |
| Set the Partition option to Dynamic range. |

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, chat or text message Description automatically generated
P1: Set the Partition option to Dynamic Range.
The SQL Server connector in copy activity provides built-in data partitioning to copy data in parallel. P2: Set the Copy method to PolyBase
Polybase is the most efficient way to move data into Azure Synapse Analytics. Use the staging blob feature to achieve high load speeds from all types of data stores, including Azure Blob storage and Data Lake Store. (Polybase supports Azure Blob storage and Azure Data Lake Store by default.)
Reference:

https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse https://docs.microsoft.com/en-us/azure/data-factory/load-azure-sql-data-warehouse

NEW QUESTION 57
- (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 61
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1. You have an application that queries DB1 to generate a sales report.
You need to see the parameter values from the last time the query was executed.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Enable Last_Query_Plan_Stats in the master database
B. Enable Lightweight_Query_Profiling in DB1
C. Enable Last_Query_Plan_Stats in DB1
D. Enable Lightweight_Query_Profiling in the master database
E. Enable PARAMETER_SNIFFING in DB1

Answer: AC

Explanation:
Last_Query_Plan_Stats allows you to enable or disable collection of the last query plan statistics (equivalent to an actual execution plan) in sys.dm_exec_query_plan_stats.
Lightweight profiling can be disabled at the database level using the LIGHTWEIGHT_QUERY_PROFILING database scoped configuration: ALTER DATABASE SCOPED CONFIGURATION SET LIGHTWEIGHT_QUERY_PROFILING = OFF;.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure

NEW QUESTION 62
- (Exam Topic 5)
You plan to develop a dataset named Purchases by using Azure Databricks. Purchases will contain the following columns:
➢ ProductID
➢ ItemPrice
➢ LineTotal
➢ Quantity
➢ StoreID
➢ Minute
➢ Month
➢ Hour
➢ Year
➢ Day
You need to store the data to support hourly incremental load pipelines that will vary for each StoreID. The solution must minimize storage costs.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
df.write
```

▼	▼
.bucketBy	("*")
.partitionBy	("StoreID", "Hour")
.range	("StoreID", "Year", "Month", "Day", "Hour")
.sortBy	("Year", "Month", "Day", "Hour" "StoreID")

```
.mode ("append")
```

▼
.csv("/Purchases")
.json("/Purchases")
.parquet("/Purchases")
.saveAsTable("/Purchases")

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Box 1: .partitionBy Example:
df.write.partitionBy("y","m","d") mode(SaveMode.Append) parquet("/data/hive/warehouse/db_name.db/" + tableName) Box 2:
("Year","Month","Day","Hour","StoreID")
Box 3: .parquet("/Purchases") Reference:
https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partiti

NEW QUESTION 67
- (Exam Topic 5)
You are designing an enterprise data warehouse in Azure Synapse Analytics that will contain a table named Customers. Customers will contain credit card information.
You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers. The solution must prevent all the salespeople from viewing or inferring the credit card information.
What should you include in the recommendation?

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

Answer: B

Explanation:
Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics support dynamic data masking. Dynamic data masking limits sensitive data exposure by masking it to non-privileged users.
The Credit card masking method exposes the last four digits of the designated fields and adds a constant string as a prefix in the form of a credit card.
Example:
XXXX-XXXX-XXXX-1234

NEW QUESTION 72
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named Appl. Appl uses db1.
You need to enable Resource Governor for a App1. The solution must meet the following requirements: App1 must be able to consume all available CPU resources.
App1 must have at least half of the available CPU resources always available.
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
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

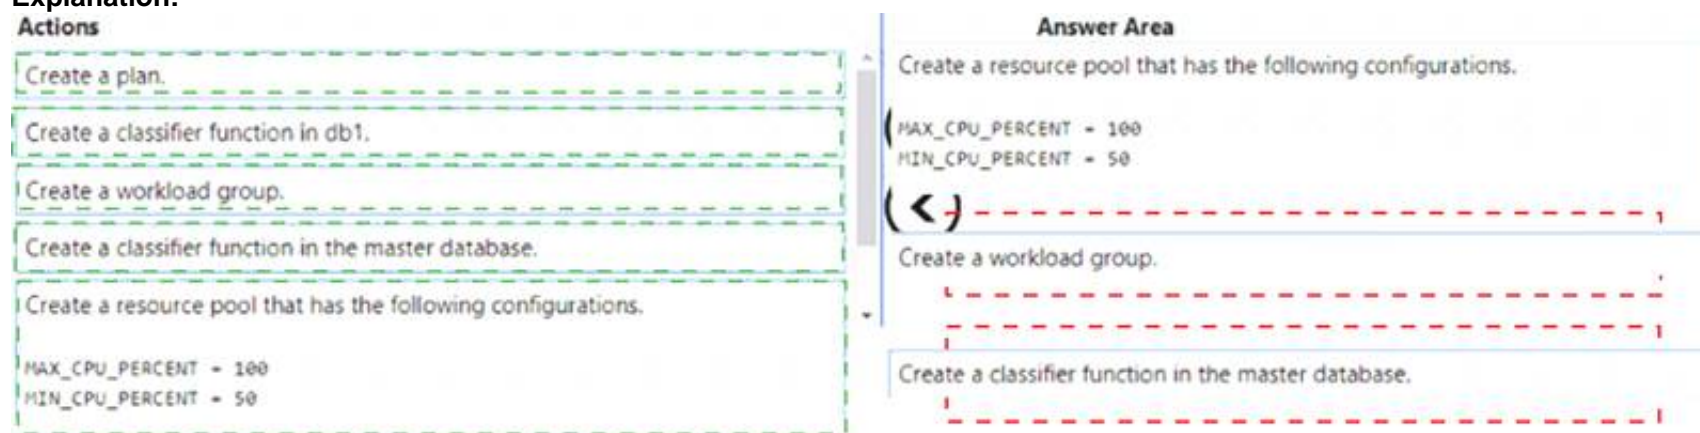| Actions	Answer Area
Create a plan.	
Create a classifier function in db1.	❯
Create a workload group.	❮
Create a classifier function in the master database.	
Create a resource pool that has the following configurations.	
MAX_CPU_PERCENT = 100 MIN_CPU_PERCENT = 50	

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
Create a plan.	Create a resource pool that has the following configurations.
Create a classifier function in db1.	MAX_CPU_PERCENT = 100 MIN_CPU_PERCENT = 50
Create a workload group.	(<)
Create a classifier function in the master database.	Create a workload group.
Create a resource pool that has the following configurations.	
MAX_CPU_PERCENT = 100 MIN_CPU_PERCENT = 50	Create a classifier function in the master database.

NEW QUESTION 74
- (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

NEW QUESTION 75
- (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 79
- (Exam Topic 5)
You have an instance of SQL Server on Azure Virtual Machine named SQL1.
You need to monitor SQL1 and query the metrics by using Kusto query language. The solution must minimize administrative effort.
Where should you store the metrics?

A. a Log Analytics workspace
B. Azure Event Hubs
C. Azure SQL Database
D. an Azure Blob storage container

Answer: A

NEW QUESTION 84
- (Exam Topic 5)
You manage an enterprise data warehouse in Azure Synapse Analytics.
Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries.
You need to monitor resource utilization to determine the source of the performance issues. Which metric should you monitor?

A. Local tempdb percentage
B. DWU percentage
C. Data Warehouse Units (DWU) used

D. Cache hit percentage

Answer: A

Explanation:
Tempdb is used to hold intermediate results during query execution. High utilization of the tempdb database can lead to slow query performance.
Note: If you have a query that is consuming a large amount of memory or have received an error message related to allocation of tempdb, it could be due to a very large CREATE TABLE AS SELECT (CTAS) or INSERT SELECT statement running that is failing in the final data movement operation.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-managemonit

NEW QUESTION 85
- (Exam Topic 5)
You have an Azure subscription that contains a server named Server1. Server1 hosts two Azure SQL databases named DB1 and DB2.
You plan to deploy a Windows app named App1 that will authenticate to DB2 by using SQL authentication. You need to ensure that App1 can access DB2. The solution must meet the following requirements:

≫ App1 must be able to view only DB2.

≫ Administrative effort must be minimized. What should you create?

A. a contained database user for App1 on DB2
B. a login for App1 on Server1
C. a contained database user from an external provider for App1 on DB2
D. a contained database user from a Windows login for App1 on DB2

Answer: D

Explanation:

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/contained-database-users-making-your-databa

NEW QUESTION 87
- (Exam Topic 5)
You have an Azure Data Factory that contains 10 pipelines.
You need to label each pipeline with its main purpose of either ingest, transform, or load. The labels must be available for grouping and filtering when using the monitoring experience in Data Factory.
What should you add to each pipeline?

A. an annotation
B. a resource tag
C. a run group ID
D. a user property
E. a correlation ID

Answer: A

Explanation:
Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag. You can define tags so you can see their performance or find errors faster.
Reference:
https://www.techtalkcorner.com/monitor-azure-data-factory-annotations/

NEW QUESTION 92
- (Exam Topic 5)
You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and an Azure Data Lake Storage Gen2 account named Account1.
You plan to access the files in Account1 by using an external table.
You need to create a data source in Pool1 that you can reference when you create the external table. 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.

```
CREATE EXTERNAL DATA SOURCE source1

WITH

    ( LOCATION = 'https://account1. [▼] .core.windows.net',
                                    ┌──────┐
                                    │ blob │
                                    │ dfs  │
                                    │ table│
                                    └──────┘

    ┌─────────────────────┐ [▼]
    │ PUSHDOWN = ON       │
    │ TYPE = BLOB_STORAGE │
    │ TYPE = HADOOP       │
    └─────────────────────┘

    )
```

A. Mastered

B. Not Mastered

Answer: A

Explanation:
Graphical user interface, table Description automatically generated
Box 1: blob
The following example creates an external data source for Azure Data Lake Gen2 CREATE EXTERNAL DATA SOURCE YellowTaxi
WITH (LOCATION = 'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/', TYPE = HADOOP)
Box 2: HADOOP
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables

NEW QUESTION 93
- (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: C

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

NEW QUESTION 98
- (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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Read onlyReadWrite

NEW QUESTION 101
- (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 schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Must use an Azure Data Factory, not an Azure Databricks job. Reference:
https://docs.microsoft.com/en-US/azure/data-factory/transform-data

NEW QUESTION 105
- (Exam Topic 5)
You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Azure region
VM1	Azure virtual machine	West US 2
MI1	Azure SQL Managed Instance	East US

You need to configure a connection between VM1 and MIL The solution must meet the following requirements:
• The connection must be encrypted.
• Network latency must be minimized. What should you implement?

A. virtual network peering
B. private endpoints
C. service endpoints
D. a site-to-site VPN

Answer: B

NEW QUESTION 106
- (Exam Topic 5)
You plan to move two 100-GB databases to Azure.
You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.
What should you use?

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

Answer: D

Explanation:
Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview

NEW QUESTION 111
- (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:
➢ Restrict network connectivity to SqlSrv1.
What should you create on VNet1?

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

Answer: C

Explanation:
A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network.
The service could be an Azure service such as:
➢ Azure Storage
➢ Azure Cosmos DB
➢ Azure SQL Database
➢ Your own service using a Private Link Service. Reference:
https://docs.microsoft.com/en-us/azure/private-link/private-endpoint-overview

NEW QUESTION 116

- (Exam Topic 5)

You have an Azure subscription.

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

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

• Minimize latency for transaction logs.

• Minimize the impact on IO Of the virtual machine.

Which type of disk should you use for each workload? To answer. drag the appropriate disk types to the correct workloads. Each disk type may be used once, at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Disk Types

| Local |
| Premium SSD |
| Standard HDD |
| Standard SSD |
| Ultra Disk |

Answer Area

TempDB: | Disk Type |

Transaction logs: | Disk Type |

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

NEW QUESTION 121

- (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 two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: From Microsoft SQL Server Management Studio (SSMS), you rename Database1 on Server2 as Database2. From the Azure portal, you create a new database on Server2 by restoring the backup of Database1 from Server1, and then you delete Database2.

Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference:

https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql

NEW QUESTION 122

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine that contains a database named Db1.

You need to enable automatic tuning for Db1.

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

NOTE: Each correct selection is worth one point.

ALTER DATABASE [Db1]

▼
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)

GO
ALTER DATABASE [Db1]

▼
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)

GO

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: SET AUTOMATIC_TUNING = AUTO
To enable automatic tuning on a single database via T-SQL, connect to the database and execute the following query:
ALTER DATABASE current SET AUTOMATIC_TUNING = AUTO
Setting automatic tuning to AUTO will apply Azure Defaults.
Box 2: SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)
To configure individual automatic tuning options via T-SQL, connect to the database and execute the query such as this one:
ALTER DATABASE current SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)
Setting the individual tuning option to ON will override any setting that database inherited and enable the tuning option. Setting it to OFF will also override any setting that database inherited and disable the tuning option.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-enable

NEW QUESTION 127
- (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 reduce the use of table variables and temporary tables. Does this meet the goal?

A. Yes
B. No

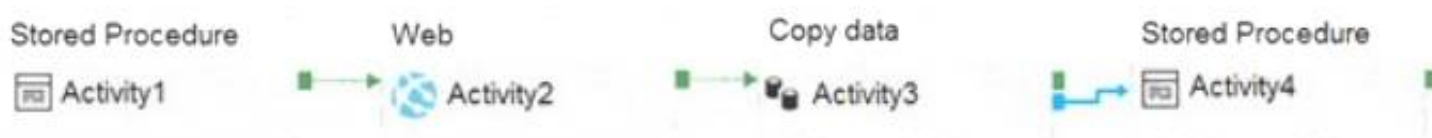**Answer:** A

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

NEW QUESTION 131
- (Exam Topic 5)
You have an Azure data factory that has two pipelines named PipelineA and PipelineB. PipelineA has four activities as shown in the following exhibit.

Stored Procedure	Web	Copy data	Stored Procedure
Activity1	Activity2	Activity3	Activity4

PipelineB has two activities as shown in the following exhibit.

Lookup	ForEach
Activity1	Activity2

You create an alert for the data factory that uses Failed pipeline runs metrics for both pipelines and all failure types. The metric has the following settings:

> Operator: Greater than
> Aggregation type: Total
> Threshold value: 2
> Aggregation granularity (Period): 5 minutes
> Frequency of evaluation: Every 5 minutes

Data Factory monitoring records the failures shown in the following table.

Pipeline	Activity	Time
PipelineA	Activity1	31-Jan-2020 10:44:00
PipelineA	Activity3	31-Jan-2020 10:47:00
PipelineB	Activity1	31-Jan-2020 10:50:00

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
An alert notification was sent after the failure of Activity1 in PipelineA.	○	○
An alert notification was sent after the failure of Activity3 in PipelineA.	○	○
An alert notification was sent after the failure of Activity1 in PipelineB.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated
Box 1: No
Just one failure within the 5-minute interval.
Box 2: No
Just two failures within the 5-minute interval.
Box 3: No
Just two failures within the 5-minute interval. Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview

NEW QUESTION 133
- (Exam Topic 5)
You have a SQL Server on Azure Virtual Machines instance that hosts a database named Db1. You need to configure the autogrow and autoshrink settings for DB1.
Which statements should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Autogrow:
- ALTER DATABASE MODIFY File and Filegroup options
- ALTER DATABASE SCOPED CONFIGURATION
- ALTER DATABASE SET options

Autoshrink:
- ALTER DATABASE MODIFY File and Filegroup options
- ALTER DATABASE SCOPED CONFIGURATION
- ALTER DATABASE SET options

A. Mastered
B. Not Mastered

Answer: A

Explanation:
https://learn.microsoft.com/en-us/troubleshoot/sql/admin/considerations-autogrow-autoshrink

NEW QUESTION 136
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL database. The database fails to respond to queries in a timely manner.

You need to identify whether the issue relates to resource_semaphore waits.
How should you complete the Transact-SQL query? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
SELECT   ┌─────────────────────┬──▼──┐
         │ is_user_process     │     │
         │ wait_time           │     │
         │ wait_type           │     │
         └─────────────────────┴─────┘

    SUM(wait_time) AS total_wait_time_ms

FROM sys.  ┌─────────────────────┬──▼──┐
           │ dm_exec_query_stats │     │
           │ dm_exec_requests    │     │
           │ query_store_query   │     │
           └─────────────────────┴─────┘

    JOIN sys.dm_exec_sessions AS dmv2
        ON dmv1.session_id = dmv2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

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/monitoring-with-dmvs

NEW QUESTION 138
- (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 140
- (Exam Topic 5)
You have an Azure SQL managed instance that hosts multiple databases.
You need to configure alerts for each database based on the diagnostics telemetry of the database. What should you use?

A. Azure SQL Analytics alerts based on metrics
B. SQL Health Check alerts based on diagnostics logs
C. SQL Health Check alerts based on metrics
D. Azure SQL Analytics alerts based on diagnostics logs

Answer: D

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

NEW QUESTION 142
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 contains a table named CustomerPII.
You need to record whenever users query the CustomerPII table.
Which two options should you enable? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. server audit specification
B. SQL Server audit
C. database audit specification
D. a server principal

Answer: AC

Explanation:
An auditing policy can be defined for a specific database or as a default server policy in Azure (which hosts SQL Database or Azure Synapse):

➢ A server policy applies to all existing and newly created databases on the server.

➢ If server auditing is enabled, it always applies to the database. The database will be audited, regardless of the database auditing settings.

➢ Enabling auditing on the database, in addition to enabling it on the server, does not override or change any of the settings of the server auditing. Both audits will exist side by side.
Note:
The Server Audit Specification object belongs to an audit.
A Database Audit Specification defines which Audit Action Groups will be audited for the specific database in which the specification is created. Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auditing-overview

NEW QUESTION 144
- (Exam Topic 5)
You have an Azure subscription that uses a domain named contoso.com.
You have two Azure VMs named DBServer1 and DBServer2. Each of them hosts a default SQL Server instance. DBServer1 is in the East US Azure region and contains a database named DatabaseA. DBServer2 is in the West US Azure region.
DBServer1 has a high volume of data changes and low latency requirements for data writes.
You need to configure a new availability group for DatabaseA. The secondary replica will reside on DBServer2.
What should you do?

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

Answer: C

Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on

NEW QUESTION 146
- (Exam Topic 5)
You have an Azure SQL managed instance.
You need to gather the last execution of a query plan and its runtime statistics. The solution must minimize the impact on currently running queries.
What should you do?

A. Generate an estimated execution plan.
B. Generate an actual execution plan.
C. Run sys.dm_exec_query_plan_scacs.
D. Generate Live Query Statistics.

Answer: C

Explanation:

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-quer

NEW QUESTION 150
- (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 executes an Azure Databricks notebook, and then inserts the data into the data warehouse.
Does this meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 155
- (Exam Topic 5)
You have an instance of SQL Server on Azure Virtual Machines that has a database named DB1. You plan to implement Azure SQL Data Sync for DB1.
Which isolation level should you configure?

A. SERIALIZABLE
B. SNAPSHOT
C. READ UNCOMMITTED
D. READ COMMITTED

Answer: B

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-data-sync-data-sql-server-sql-database

NEW QUESTION 156
- (Exam Topic 2)
You need to implement a solution to notify the administrators. The solution must meet the monitoring requirements.
What should you do?

A. Create an Azure Monitor alert rule that has a static threshold and assign the alert rule to an action group.
B. Add a diagnostic setting that logs QueryStoreRuntimeStatistics and streams to an Azure event hub.
C. Add a diagnostic setting that logs Timeouts and streams to an Azure event hub.
D. Create an Azure Monitor alert rule that has a dynamic threshold and assign the alert rule to an action group.

Answer: D

Explanation:
Reference:
https://azure.microsoft.com/en-gb/blog/announcing-azure-monitor-aiops-alerts-with-dynamic-thresholds/

NEW QUESTION 158
- (Exam Topic 2)
You are evaluating the role assignments.
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
DBAGroup1 will be able to sign in to each customer's Azure SQL database by using Azure Data Studio.	○	○
DBAGroup1 will be able to assign the SQL DB Contributor role to other users.	○	○
DBAGroup2 will be able to create a new Azure SQL database on each customer's Azure SQL Database server.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
DBAGroup1 is member of the Contributor role.
The Contributor role grants full access to manage all resources, but does not allow you to assign roles in Azure RBAC, manage assignments in Azure Blueprints, or share image galleries.
Box 2: No
Box 3: Yes
DBAGroup2 is member of the SQL DB Contributor role.
The SQL DB Contributor role lets you manage SQL databases, but not access to them. Also, you can't manage their security-related policies or their parent SQL servers. As a member of this role you can create and manage SQL databases.
Reference:
https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles

NEW QUESTION 162
- (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 165
- (Exam Topic 1)
You need to implement authentication for ResearchDB1. The solution must meet the security and compliance requirements.
What should you run as part of the implementation?

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

Answer: E

Explanation:
Scenario: Authenticate database users by using Active Directory credentials.
(Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.) Authenticate the user in SQL Database or SQL Data Warehouse based on an Azure Active Directory user: CREATE USER [Fritz@contoso.com] FROM EXTERNAL PROVIDER;
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql

NEW QUESTION 166
- (Exam Topic 1)
You are planning the migration of the SERVER1 databases. The solution must meet the business requirements.
What should you include in the migration plan? 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:
Azure Database Migration service
Box 1: Premium 4-VCore
Scenario: Migrate the SERVER1 databases to the Azure SQL Database platform.
Minimize downtime during the migration of the SERVER1 databases.

Premimum 4-vCore is for large or business critical workloads. It supports online migrations, offline migrations, and faster migration speeds.
Reference: https://azure.microsoft.com/pricing/details/database-migration/
https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-azure-sql-online

NEW QUESTION 170
- (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.

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 172
- (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 177
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

DP-300 Practice Exam Features:

* DP-300 Questions and Answers Updated Frequently

* DP-300 Practice Questions Verified by Expert Senior Certified Staff

* DP-300 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DP-300 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DP-300 Practice Test Here](https://www.certshared.com/exam/DP-300/)