

DP-300^{Q&As}

Administering Relational Databases on Microsoft Azure

Pass Microsoft DP-300 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.certbus.com/dp-300.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

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

Correct Answer: A

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>

QUESTION 2

DRAG DROP

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.

Select and Place:

Actions

Answer Area

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.

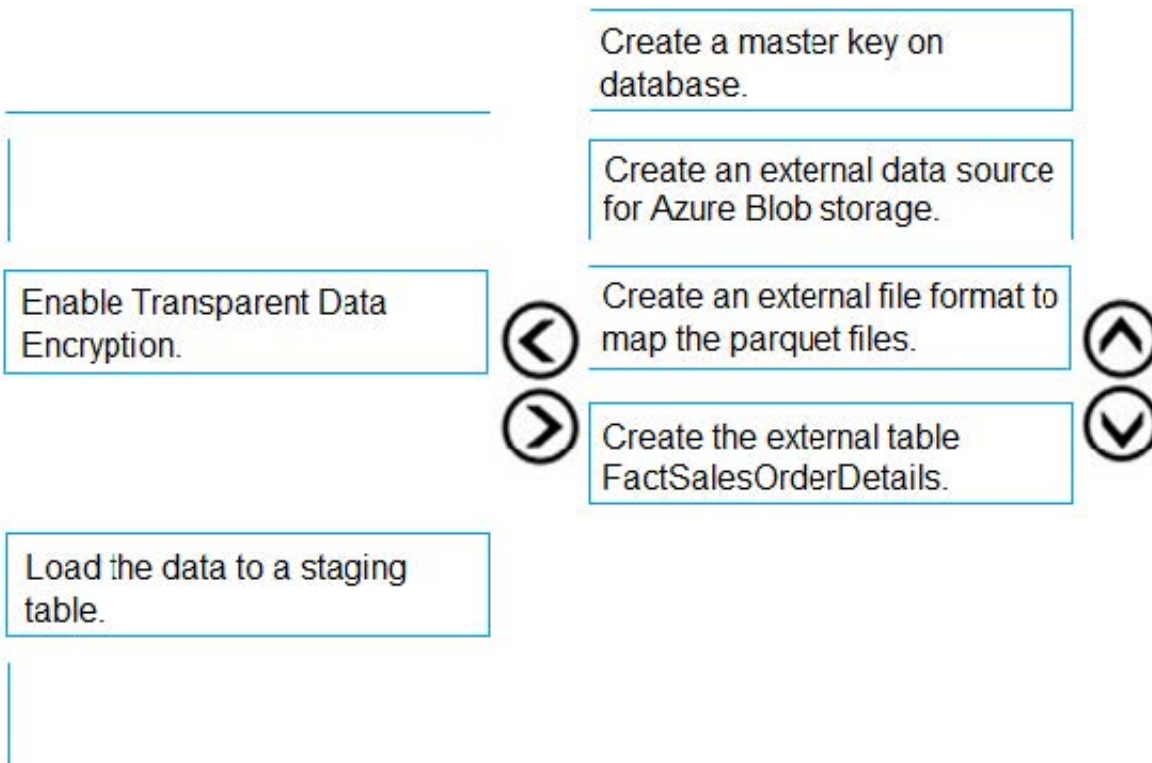
Create an external file format to map the parquet files.



Correct Answer:

Actions

Answer Area



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>

QUESTION 3

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

Correct Answer: C

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>

QUESTION 4

You have An Azure SQL managed instance.

You need to configure the SQL Server Agent service to email job notifications.

Which statement should you execute?

- A. EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sysadmin_dbmail_profile';
- B. EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'application_dbmail_profile';
- C. EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'AzureManagedInstance_dbmail_profile';
- D. EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sys_dbmail_profile';

Correct Answer: C

-- Create a Database Mail profile

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = '\\AzureManagedInstance_dbmail_profile\\',  
@description = '\\E-mail profile used for messages sent by Managed Instance SQL Agent.\\';
```

Reference: <https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/job-automation-managed-instance>

QUESTION 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. An Azure SQL Database elastic pool
- B. SQL Server on Azure virtual machines
- C. an Azure SQL Database managed instance
- D. Azure SQL databases

Correct Answer: A

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price.

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

QUESTION 6

DRAG DROP

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.

Select and Place:

Actions

- 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.

Answer Area



Correct Answer:

Actions

- Deploy Grafana to an Azure virtual machine.
-
-
- Create a data source in Azure Monitor.
-

Answer Area

- Configure the Databricks cluster to use the Databricks monitoring library.
- Build a **spark-listeners-loganalytics-1.0-SNAPSHOT.jar** JAR file.
- Create Dropwizard counters in the application code.



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

QUESTION 7

DRAG DROP

You need to configure user authentication for the SERVER1 databases. 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.

Select and Place:

Actions

Answer Area

Modify the Azure SQL server administrator account.

Create database users.

Create a user in the master database.

Create an Azure AD administrator for the logical server.

Connect to the databases by using an Azure AD account.



Correct Answer:

Actions

Modify the Azure SQL server administrator account.

Create a user in the master database.

Answer Area

Create an Azure AD administrator for the logical server.

Create database users.

Connect to the databases by using an Azure AD account.



Scenario: Authenticate database users by using Active Directory credentials.

The configuration steps include the following procedures to configure and use Azure Active Directory authentication.

1.
Create and populate Azure AD.
2.
Optional: Associate or change the active directory that is currently associated with your Azure Subscription.
3.
Create an Azure Active Directory administrator. (Step 1)
4.
Configure your client computers.
5.
Create contained database users in your database mapped to Azure AD identities. (Step 2)
6.
Connect to your database by using Azure AD identities. (Step 3)

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

QUESTION 8

HOTSPOT

You have the following Azure Resource Manager template.

```

...
  "variable": {
    "serverName": "azsqlserver0001"
  },
  "resources": [
    {
      "name": "[variables('serverName')]",
      "type": "Microsoft.Sql/servers",
      "apiVersion": "2019-06-01-preview",
      "location": "[parameters('location')]",
      "properties": {
        "administratorLogin": "[parameters('administratorLogin')]",
        "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
        "version": "12.0"
      },
      "resources": [
        {
          "name": "[concat(variables('serverName'),'/',parameters('databaseName'))]",
          "type": "Microsoft.Sql/servers/databases",
          "apiVersion": "2020-08-01-preview",
          "location": "[parameters('location')]",
          "kind": "v12.0"
          "sku": {
            "name": "Standard",
            "tier": "Standard",
            "capacity": 10
          },
          "dependsOn": [
            "[concat('Microsoft.Sql/servers/', variables('serverName'))]"
          ],
          "properties": {
          },
          "resources": [
          ]
        }
      ]
    }
  ],
  ...

```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
The template deploys a serverless Azure SQL database.	<input type="radio"/>	<input type="radio"/>
The template deploys a database to an Azure SQL Database managed instance.	<input type="radio"/>	<input type="radio"/>
The pricing tier of the database deployment is based on DTUs.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Statements	Yes	No
The template deploys a serverless Azure SQL database.	<input type="radio"/>	<input checked="" type="radio"/>
The template deploys a database to an Azure SQL Database managed instance.	<input type="radio"/>	<input checked="" type="radio"/>
The pricing tier of the database deployment is based on DTUs.	<input checked="" type="radio"/>	<input type="radio"/>

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

QUESTION 9

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

Correct Answer: D

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

QUESTION 10

HOTSPOT

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:

1.

Support user-initiated backups.

2.

Support multiple automatically replicated instances across Azure regions.

3.

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.

Hot Area:

Answer Area

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	

Correct Answer:

Answer Area

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	

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-is-overview>

<https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure>

QUESTION 11

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

You need to retrieve query plans from the Query Store on DB1.

What should you do first?

- A. From Microsoft SQL Server Management Studio, modify the properties of DB1.
- B. From Microsoft SQL Server Management Studio, modify the properties of the SQL Server instance.
- C. On SQLVM1, install the SQL Server IaaS Agent extension.
- D. On SQLVM1, install the Azure Monitor agent for Windows.

Correct Answer: A

Enable the Query Store Use the Query Store page in SQL Server Management Studio

1.

In Object Explorer, right-click a database, and then select Properties.

2.

In the Database Properties dialog box, select the Query Store page.

3.

In the Operation Mode (Requested) box, select Read Write.

Reference: <https://learn.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store>

QUESTION 12

You have the following resources:

1.

15 SQL Server on Azure Virtual Machines instances

2.

20 Azure SQL databases

You need to recommend a solution to centrally monitor the resources for security vulnerabilities.

What should you include in the recommendation?

- A. database audits
- B. Microsoft Defender
- C. SQL insights
- D. Azure SQL Auditing

Correct Answer: B

Keyword is the "vulnerabilities" <https://learn.microsoft.com/en-us/azure/defender-for-cloud/defender-for-sql-introduction>

QUESTION 13

You have an Azure SQL database named DB1.

You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails. The solution must minimize costs.

Which deployment option and pricing tier should you configure?

- A. Azure SQL Database Business Critical
- B. Azure SQL Database Standard
- C. Azure SQL Database managed instance General Purpose
- D. Azure SQL Database Hyperscale

Correct Answer: A

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). The routing to a specific gateway ring is controlled by Azure Traffic Manager

(ATM). Because the zone redundant configuration in the Premium or Business Critical service tiers does not create additional database redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration, you can

make your Premium or Business Critical databases resilient to a much larger set of failures, including catastrophic datacenter outages, without any changes to the application logic. You can also convert any existing Premium or

Business

Critical databases or pools to the zone redundant configuration.

Incorrect Answers:

Note:

There are several versions of this question with different incorrect answer options. Other incorrect answers you may see on the exam include:

1.

Azure SQL Database Hyperscale

2.

Azure SQL Database Basic

3.

Azure SQL Database managed instance General Purpose

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

QUESTION 14

You have an Azure virtual machine named VM1 that runs Windows Server 2022 and hosts a Microsoft SQL Server 2019 instance named SQL1.

You need to configure SQL1 to use mixed mode authentication.

Which procedure should you run?

A. sp_addremotelogin

B. xp_instance_regwrite

C. sp_change_users_login

D. xp_grant_login

Correct Answer: B

Change authentication mode (Transact-SQL)

The following example changes Server Authentication from mixed mode (Windows and SQL) to Windows only.

```
USE [master]
```

```
GO
```

```
EXEC xp_instance_regwrite N\\HKEY_LOCAL_MACHINE\\,
```

```
N\\Software\\Microsoft\\MSSQLServer\\MSSQLServer\\, N\\LoginMode\\, REG_DWORD, 1; GO
```


Reference: <https://learn.microsoft.com/en-us/sql/database-engine/configure-windows/change-server-authentication-mode>

QUESTION 15

You are designing a dimension table in an Azure Synapse Analytics dedicated SQL pool.

You need to create a surrogate key for the table. The solution must provide the fastest query performance.

What should you use for the surrogate key?

- A. an IDENTITY column
- B. a GUID column
- C. a sequence object

Correct Answer: A

Dedicated SQL pool supports many, but not all, of the table features offered by other databases. Surrogate keys are not supported. Implement it with an Identity column.

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

[DP-300 Practice Test](#)

[DP-300 Study Guide](#)

[DP-300 Exam Questions](#)