

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.leads4pass.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 have an instance of SQL Server on Azure Virtual Machines.

You need to ensure that a user named User1 can configure proxy accounts for SQL Server Agent jobs. The solution must use the principle of least privilege.

Which role should you assign to User1?

- A. sysadmin
- B. SQLAgentUserRole
- C. SQLAgentReaderRole
- D. SQLAgentOperatorRole

Correct Answer: A

Permissions

Only members of the sysadmin fixed server role have permission to create, modify, or delete proxy accounts. Users who are not members of the sysadmin fixed server role must be added to one of the following SQL Server Agent fixed

database roles in the msdb database to use proxies: SQLAgentUserRole, SQLAgentReaderRole, or SQLAgentOperatorRole.

Incorrect:

Not B: SQLAgentUserRole is the least privileged of the SQL Server Agent fixed database roles. It has permissions on only operators, local jobs, and job schedules. Members of SQLAgentUserRole have permissions on only local jobs and job

schedules that they own.

Not C: SQLAgentReaderRole includes all the SQLAgentUserRole permissions as well as permissions to view the list of available multiserver jobs, their properties, and their history. Members of this role can also view the list of all available jobs

and job schedules and their properties, not just those jobs and job schedules that they own.

Not D: SQLAgentOperatorRole is the most privileged of the SQL Server Agent fixed database roles. It includes all the permissions of SQLAgentUserRole and SQLAgentReaderRole. Members of this role can also view properties for operators

and proxies, and enumerate available proxies and alerts on the server.

Reference:

<https://learn.microsoft.com/en-us/sql/ssms/agent/create-a-sql-server-agent-proxy>

<https://learn.microsoft.com/en-us/sql/ssms/agent/sql-server-agent-fixed-database-roles>

QUESTION 2

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

Correct Answer: B

If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity, not an Azure Databricks notebook, 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>

QUESTION 3**HOTSPOT**

You have an instance of SQL Server on Azure Virtual Machines named VM1.

You need to use an Azure Automation runbook to initiate a SQL Server database backup on VM1.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

-Name "BackupDB" -ResourceGroupName "RG1"

Move-AzAutomationHybridRunbookWorker
Search-ADAccount
Set-AzAutomationRunBook
Start-AzAutomationRunBook

"Backups" -Wait

-AutomationAccountName
-Parameters
-RunOn

Correct Answer:

Answer Area

-Name "BackupDB" -ResourceGroupName "RG1"

Move-AzAutomationHybridRunbookWorker
Search-ADAccount
Set-AzAutomationRunBook
Start-AzAutomationRunBook

"Backups" -Wait

-AutomationAccountName
-Parameters
-RunOn

Explanation:

Box 1: Start-AzAutomationRunBook

The Start-AzAutomationRunbook cmdlet starts an Azure Automation runbook job. Specify the ID or name of a runbook.

Syntax:

```
Start-AzAutomationRunbook [-Name] [-Parameters] [-RunOn] [-Wait] [-MaxWaitSeconds] [-ResourceGroupName] [-AutomationAccountName] [-DefaultProfile] []
```

Example: Start a runbook job

```
Start-AzAutomationRunbook -AutomationAccountName "Contoso17" -Name "Runbk01" -ResourceGroupName
```

"ResourceGroup01"

This command starts a runbook job for the runbook named Runbk01 in the Azure Automation account named Contoso17.

Incorrect:

* ADAccount Box 2: -AutomationAccountName Reference:

<https://learn.microsoft.com/en-us/powershell/module/az.automation/start-azautomationrunbook>

QUESTION 4

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:

1.

Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.

2.

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

Correct Answer: C

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>

QUESTION 5

DRAG DROP 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? To answer. move the appropriate actions from the list of actions to answer

area and arrange them in correct order.

Select and Place:

Actions

Create a Database Mail account.

Enable pager notifications upon failure.

Create a profile named `AzureManagedInstance_dbmail_profile`.

Enable email notifications upon failure.

Create a profile named `application_dbmail_profile`.

Answer Area

Correct Answer:

Actions

Enable pager notifications upon failure.
Create a profile named <code>application_dbmail_profile</code> .

Answer Area

Create a Database Mail account.
Create a profile named <code>AzureManagedInstance_dbmail_profile</code> .
Enable email notifications upon failure.

QUESTION 6

DRAG DROP

You create an Azure SQL managed instance and a job that performs backups.

You need to configure the job to notify a distribution group by email when the job fails. 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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:

Actions Commands Cmdlets Statements

- Configure Database Mail.
- Configure a job notification.
- Configure an alert.
- Configure SendGrid.
- Create an operator.

Answer Area

Correct Answer:

Actions Commands Cmdlets Statements

-
-
- Configure an alert.
- Configure SendGrid.
-

Answer Area

Automate management tasks using SQL Agent jobs in Azure SQL Managed Instance

Using SQL Server Agent in SQL Server and SQL Managed Instance, you can create and schedule jobs that could be periodically executed against one or many databases to run Transact-SQL (T-SQL) queries and perform maintenance tasks.

Step 1: Configure Database Mail

If it isn't already enabled, first you would need to configure the Database Mail feature on SQL Managed Instance:

GO

```
EXEC sp_configure 'show advanced options', 1;
```

GO


```
RECONFIGURE;
```

```
GO
```

```
EXEC sp_configure '\\Database Mail XPs\\', 1;
```

```
GO
```

```
RECONFIGURE
```

Step 2: Create an operator

You can notify the operator that something happened with your SQL Agent jobs. An operator defines contact information for an individual responsible for the maintenance of one or more instances in SQL Managed Instance. Sometimes, operator responsibilities are assigned to one individual.

You can create operators using SQL Server Management Studio (SSMS) or the Transact-SQL script shown in the following example:

```
EXEC msdb.dbo.sp_add_operator @name=N'AzureSQLTeam', @enabled=1,  
@email_address=N'AzureSQLTeamn@contoso.com';
```

Step 3: Configure a job notification

Job notifications

SQL Agent jobs enable you to get notifications when the job finishes successfully or fails. You can receive notifications via email.

Example:

You can then modify any SQL Agent job and assign operators that will be notified via email if the job completes, fails, or succeeds using SSMS or the following T-SQL script:

```
EXEC msdb.dbo.sp_update_job @job_name=N'Load data using SSIS', @notify_level_email=3, -- Options are: 1 on  
succeed, 2 on failure, 3 on complete @notify_email_operator_name=N'AzureSQLTeam';
```

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

QUESTION 7

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

Correct Answer: C

The lifecycle management policy lets you:

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

Scenario:

1. Purge Twitter feed data records that are older than two years.
2. Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.
3. Minimize administrative effort to maintain the Twitter feed data records.

Incorrect Answers:

A: Time-based retention policy support: Users can set policies to store data for a specified interval. When a time-based retention policy is set, blobs can be created and read, but not modified or deleted. After the retention period has expired, blobs can be deleted but not overwritten.

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

QUESTION 8

DRAG DROP

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.

Select and Place:

Actions

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

Answer Area



Correct Answer:

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

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>

QUESTION 9

HOTSPOT

You have an Azure data factory that has two pipelines named PipelineA and PipelineB.

PipelineA has four activities as shown in the following exhibit.



PipelineB has two activities as shown in the following exhibit.



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:

1.

Operator: Greater than

2.

Aggregation type: Total

3.

Threshold value: 2

4.

Aggregation granularity (Period): 5 minutes

5.

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.

Hot Area:

Answer Area

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity3 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity1 in PipelineB.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	<input type="radio"/>	<input checked="" type="radio"/>
An alert notification was sent after the failure of Activity3 in PipelineA.	<input type="radio"/>	<input checked="" type="radio"/>
An alert notification was sent after the failure of Activity1 in PipelineB.	<input type="radio"/>	<input checked="" type="radio"/>

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>

QUESTION 10

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 11

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

Correct Answer: B

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>

QUESTION 12

You have an Always On availability group deployed to Azure virtual machines. The availability group contains a

database named DB1 and has two nodes named SQL1 and SQL2. SQL1 is the primary replica.

You need to initiate a full backup of DB1 on SQL2.

Which statement should you run?

- A. `BACKUP DATABASE DB1 TO URL=\\https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak\\ with (Differential, STATS=5, COMPRESSION);`
- B. `BACKUP DATABASE DB1 TO URL=\\https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak\\ with (COPY_ONLY, STATS=5, COMPRESSION);`
- C. `BACKUP DATABASE DB1 TO URL=\\https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak\\ with (File_Snapshot, STATS=5, COMPRESSION);`
- D. `BACKUP DATABASE DB1 TO URL=\\https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak\\ with (Nolnit, STATS=5, COMPRESSION);`

Correct Answer: B

BACKUP DATABASE supports only copy-only full backups of databases, files, or filegroups when it's executed on secondary replicas. Copy-only backups don't impact the log chain or clear the differential bitmap. Incorrect Answers:

A: Differential backups are not supported on secondary replicas. The software displays this error because the secondary replicas support copy-only database backups.

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/active-secondaries-backup-on-secondary-replicas-always-on-availability-groups>

QUESTION 13

You have an Azure data solution that contains an enterprise data warehouse in Azure Synapse Analytics named DW1.

Several users execute adhoc queries to DW1 concurrently.

You regularly perform automated data loads to DW1.

You need to ensure that the automated data loads have enough memory available to complete quickly and successfully when the adhoc queries run.

What should you do?

- A. Assign a smaller resource class to the automated data load queries.
- B. Create sampled statistics to every column in each table of DW1.
- C. Assign a larger resource class to the automated data load queries.
- D. Hash distribute the large fact tables in DW1 before performing the automated data loads.

Correct Answer: C

The performance capacity of a query is determined by the user's resource class. Smaller resource classes reduce the

maximum memory per query, but increase concurrency. Larger resource classes increase the maximum memory per query, but reduce concurrency.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/resource-classes-for-workload-management>

QUESTION 14

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

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

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

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

What should you do?

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

Correct Answer: A

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

QUESTION 15

HOTSPOT

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.

Hot Area:

Answer Area

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	

Correct Answer:

Answer Area

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	

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.

Incorrect:

SQL authentication: To connect to dedicated SQL pool (formerly SQL DW), you must provide the following information:

1.

Fully qualified servername

2.

Specify SQL authentication

3.

Username

4.

Password

5.

Default database (optional)

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

[Latest DP-300 Dumps](#)

[DP-300 VCE Dumps](#)

[DP-300 Study Guide](#)