



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.pass4itsure.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 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

Correct Answer: D

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-export-configure?tabs=azure-portal#configure-the-streaming-export-of-diagnostic-telemetry>

QUESTION 2

DRAG DROP

You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1. DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL

Server Data Tools (SSDT) database project.

You need to populate a new resource group named App1Test with the DB1 database and an Azure SQL Server named TestServer1. The resources in App1Test must have the same configurations as the resources in App1Dev.

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

Change the Active Directory Admin on TestServer1

Change the server name and related variables in the templates

From the database project, deploy the database schema and permissions

Add IP addresses to the firewall

From the Azure portal, export the Azure Resource Manager templates

From the Azure portal, deploy the templates.



Correct Answer:

Actions

Answer Area

Change the Active Directory Admin on TestServer1

Add IP addresses to the firewall

From the Azure portal, export the Azure Resource Manager templates

Change the server name and related variables in the templates

From the Azure portal, deploy the templates.

From the database project, deploy the database schema and permissions



QUESTION 3

You have an Azure subscription.

You create a logical SQL server that hosts four databases. Each database will be used by a separate customer.

You need to ensure that each customer can access only its own database. The solution must minimize administrative effort.



Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Deny public access.
- B. Create a private endpoint.
- C. Create a database-level firewall rule.
- D. Create a network security group (NSG).
- E. Create a server-level firewall rule.

Correct Answer: B

You can connect to an Azure SQL server using an Azure Private Endpoint.

Azure Private endpoint is the fundamental building block for Private Link in Azure. It enables Azure resources, like virtual machines (VMs), to privately and securely communicate with Private Link resources such as Azure SQL server.

Reference:

<https://learn.microsoft.com/en-us/azure/private-link/tutorial-private-endpoint-sql-portal>

QUESTION 4

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 mapping data flow, 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 a mapping flow, 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 5

HOTSPOT

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

Type	Designated retention period
Application	360 days
Infrastructure	60 days

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

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

1. Automatically deletes the logs at the end of each retention period
2. Minimizes storage costs

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

To minimize storage costs:

<input type="checkbox"/>	Store the infrastructure logs and the application logs in the Archive access tier.
<input type="checkbox"/>	Store the infrastructure logs and the application logs in the Cool access tier.
<input type="checkbox"/>	Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

<input type="checkbox"/>	Azure Data Factory pipelines
<input type="checkbox"/>	Azure Blob storage lifecycle management rules
<input type="checkbox"/>	Immutable Azure Blob storage time-based retention policies

Correct Answer:



Answer Area

To minimize storage costs:

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

To delete the logs automatically:

- Azure Data Factory pipelines
- Azure Blob storage lifecycle management rules
- Immutable Azure Blob storage time-based retention policies

Box 1: Store the infrastructure logs in the Cool access tier the application logs in the Archive access tier

Hot - Optimized for storing data that is accessed frequently.

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

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

Box 2: Azure Blob storage lifecycle management rules

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

Reference:

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

QUESTION 6

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 VM1 cannot connect to any Azure SQL Server other than SqlSrv1.

2.

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



Correct Answer: C

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:

1.

Azure Storage

2.

Azure Cosmos DB

3.

Azure SQL Database


4.

Your own service using a Private Link Service.

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

QUESTION 7

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



You need to reduce the Log IO percentage. The solution must minimize costs.

What should you do?

- A. Change Service tier to Business Critical.
- B. Increase the number of vCores.
- C. Perform a checkpoint operation.
- D. Change Recovery model to Simple.

Correct Answer: A

High Log IO is usually seen when you hit the IOPS/throughput limits for the respective pricing tier.

* Data/Log IOPS (approximate) General Purpose service tier: 500 - 7500 per file Increase file size to get more IOPS

Business Critical:

16 K - 320 K (4000 IOPS/vCore)

Add more vCores to get better IO performance.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/resource-limits>

QUESTION 8



HOTSPOT

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:

1.
The first six numerals of each customer's phone number must be masked.
 2.
The last four digits of each customer's phone number must be visible.
 3.
Hyphens must be preserved and displayed.
- How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

Hot Area:



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	

Correct Answer:



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	

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>

QUESTION 9

HOTSPOT

You have an Azure subscription that contains a resource group named RG1. RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL

You need to use PowerShell to enable and configure automated patching for SQL The solution must include both SQL Server and Windows security updates.

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

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

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

```
Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |
```

	▼
Get-AzVMSSQLServerExtension	
Set-AzVMExtension	
Set-AzVMSqlServerExtension	

```
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
```

```
-SQLManagementType
```

	▼
-Full	
-Lightweight	
-NoAgent	


Correct Answer:



Answer Area

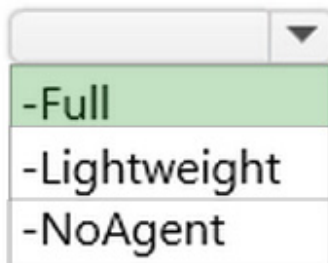
```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable  
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2  
-MaintenanceWindowDuration 120 -PatchCategory "Important"
```

```
Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |
```



```
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
```

```
-SQLManagementType
```



Box 1: Set-AzVMSqlServerExtension

The Set-AzVMSqlServerExtension cmdlet sets the AzureSQL Server extension on a virtual machine.

Example: Set automatic patching settings on a virtual machine.

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable -DayOfWeek "Thursday"  
-MaintenanceWindowStartingHour 11 -MaintenanceWindowDuration 120 -PatchCategory "Important"
```

```
Get-AzVM -ResourceGroupName "testrg" -Name "VirtualMachine11" | Set-AzVMSqlServerExtension  
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
```

The first command creates a configuration object by using the New-AzVMSqlServerAutoPatchingConfig cmdlet. The command stores the configuration in the \$AutoPatchingConfig variable. The second command gets the virtual machine

named VirtualMachine11 in the Resource Group testrg by using the Get-AzVM cmdlet. The command passes that object to the current cmdlet by using the pipeline operator. The current cmdlet sets the automatic patching settings in



\$AutoPatchingConfig for the virtual machine. The command passes the virtual machine to the Update-AzVM cmdlet.

Box 2: -Full

The solution must include both SQL Server and Windows security updates.

Reference:

<https://learn.microsoft.com/en-us/powershell/module/az.compute/set-azvmsqlserverextension>

QUESTION 10

HOTSPOT

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.

Hot Area:



```
"variables": (  
  "dataDisks": {  
    "caching":  "dataDiskCount": 8,"logDisksCount": 1,  
    ...  
  }  
)  
"resources": [  
  ...  
  { "type": "Microsoft.Compute/virtualMachines",  
    ...  
    "storageProfile": {  
      "osDisk": {  
        ...  
      }, "copy": [  
        {"name": "dataDisks","count": "[add(variables('dataDiskCount'),  
          variables('logDisksCount'))]"},  
        "input": {"lun": "[copyIndex('dataDisks')]", "createOption": "empty",  
        "caching": "[if(greaterOrEquals(copyIndex('dataDisks'),  
          parameters('dataDiskCount')),  
           variables('dataDisks').caching )]", "diskSizeGB": 1023,  
        ...  
      ]  
    }  
  ]  
}
```

Correct Answer:



```
"variables": (  
  "dataDisks": {  
    "caching":  "dataDiskCount": 8,"logDisksCount": 1,  
    ...  
  }  
)  
"resources": [  
  ...  
  { "type": "Microsoft.Compute/virtualMachines",  
    ...  
    "storageProfile": {  
      "osDisk": {  
        ...  
      }, "copy": [  
        {"name": "dataDisks","count": "[add(variables('dataDiskCount'),  
          variables('logDisksCount'))]"},  
        "input": {"lun": "[copyIndex('dataDisks')]", "createOption": "empty",  
        "caching": "[if(greaterOrEquals(copyIndex('dataDisks'),  
          parameters('dataDiskCount')),  
           variables('dataDisks').caching )]", "diskSizeGB": 1023,  
        ...  
      ]
```

Box 1: None

Set host caching to none for log file disks.

Box 2: ReadOnly



Set host caching to read-only for data file disks.

Do not enable read/write caching on disks that contain SQL Server data or log files.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/performance-guidelines-best-practices-storage>

QUESTION 11

You have an Azure subscription that contains two instances of SQL Server on Azure Virtual Machines named VM1 and VM2. Both instances run Microsoft SQL Server 2019 CU8.

You need to deploy a failover cluster instance (FCI) to VM1 and VM2 that will use Azure shared disks. The solution must maximize resiliency.

Which quorum option should you use?

- A. node majority with a cloud witness
- B. node majority with no witness
- C. node majority with a file share witness
- D. node majority with a disk witness

Correct Answer: D

Configure quorum

Since the disk witness is the most resilient quorum option, and the FCI solution uses Azure shared disks, it's recommended to configure a disk witness as the quorum solution.

Reference: <https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/failover-cluster-instance-azure-shared-disks-manually-configure>

QUESTION 12

HOTSPOT

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

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

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

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

Hot Area:



Answer Area

To prepare for the migration:

Add a secondary replica to AG1.
Create a SQL Server 2019 Always On availability group on VM1.
Upgrade the on-premises SQL servers to SQL Server 2019.

To perform the migration, use:

A distributed availability group
Azure Migrate
Log shipping

Correct Answer:

Answer Area

To prepare for the migration:

Add a secondary replica to AG1.
Create a SQL Server 2019 Always On availability group on VM1.
Upgrade the on-premises SQL servers to SQL Server 2019.

To perform the migration, use:

A distributed availability group
Azure Migrate
Log shipping

Box 1: Add a secondary replica to AG1.

Migrate an availability group to SQL Server on Azure VM

To ensure the migrated server is synchronized with the source server, stop the SQL Server service on every replica in the availability group, starting with secondary replicas (in SQL Server Configuration Manager > Services) while ensuring



the disks hosting SQL data are online.

Incorrect:

*

Upgrade to SQL Server 2019

No need to upgrade to SQL Server 2019

*

Create a SQL Server 2019 Always ON available group on VM1.

You do that afterwards.

After the migration, enable Always On availability groups for the database.

Box 2: Azure Migrate You can migrate your SQL Server Always On availability group to SQL Server on Azure VMs using the Azure Migrate: Server Migration tool. Using the migration tool, you will be able to migrate each replica in the availability group to an Azure VM hosting SQL Server, as well as the cluster metadata, availability group metadata and other necessary high availability components.

Reference: <https://learn.microsoft.com/en-us/azure/azure-sql/migration-guides/virtual-machines/sql-server-availability-group-to-sql-on-azure-vm>

QUESTION 13

HOTSPOT

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.

Hot Area:



Answer Area

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 dmvs

ON dmvs1.session_id = dmvs2.session_id

WHERE is_user_process = 1

GROUP BY wait_type

ORDER BY SUM(wait_time) DESC;

Correct Answer:



Answer Area

```
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 dmvs2
```

```
        ON dmvs1.session_id = dmvs2.session_id
```

```
WHERE is_user_process = 1
```

```
GROUP BY wait_type
```

```
ORDER BY SUM(wait_time) DESC;
```

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs>

[DP-300 PDF Dumps](#)

[DP-300 Practice Test](#)

[DP-300 Study Guide](#)