



Actual exam question from Microsoft's DP-300

Question #: 1

Topic #: 1

[\[All DP-300 Questions\]](#)

You have 20 Azure SQL databases provisioned by using the vCore purchasing model.

You plan to create an Azure SQL Database elastic pool and add the 20 databases.

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

NOTE: Each correct selection is worth one point.

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 2

Topic #: 1

[\[All DP-300 Questions\]](#)

DRAG DROP -

You have SQL Server 2019 on an Azure virtual machine that contains an SSISDB database.

A recent failure causes the master database to be lost.

You discover that all Microsoft SQL Server integration Services (SSIS) packages fail to run on the virtual machine.

Which four actions should you perform in sequence to resolve the issue? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct.

Select and Place:

Actions

Answer Area

Add a certificate to an Azure key vault

Enable Transparent Data Encryption (TDE)

Encrypt a copy of the master key by using the service master key

Turn on the TRUSTWORTHY property and the CLR property

Attach the SSISDB database

Open the master key for the SSISDB database



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 3

Topic #: 1

[\[All DP-300 Questions\]](#)

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. You must provide the greatest reduction in space for the database and maximize performance.

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

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 4

Topic #: 1

[\[All DP-300 Questions\]](#)

You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

- ⇒ Clustered columnstore indexes
- ⇒ Automatic tuning
- ⇒ Change tracking
- ⇒ PolyBase

You plan to migrate DB1 to an Azure SQL database.

What feature should be removed or replaced before DB1 can be migrated?

- A. Clustered columnstore indexes
- B. PolyBase
- C. Change tracking
- D. Automatic tuning

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 5

Topic #: 1

[\[All DP-300 Questions\]](#)

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. Always On Availability Group
- D. Azure Database Migration Service

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 6

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

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.

Hot Area:

Answer Area

Backup type:

	▼
Full and log backups only	
Full backup only	
Log backup only	

Backup option:

	▼
WITH CHECKSUM	
WITH NOINIT	
WITH UNLOAD	

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 7

Topic #: 1

[\[All DP-300 Questions\]](#)

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

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.

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 8

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

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.

Hot Area:

Answer Area

```
CREATE EXTERNAL DATA SOURCE source1
```

```
WITH
```

```
( LOCATION = 'https://account1.  .core.windows.net',
```

<input type="text"/>
blob
dfs
table

<input type="text"/>
PUSHDOWN = ON
TYPE = BLOB_STORAGE
TYPE = HADOOP

```
)
```

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 9

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

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.

Hot Area:

Answer Area

```
df.write
```

<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> .bucketBy	<input type="checkbox"/> ("*")
<input type="checkbox"/> .partitionBy	<input type="checkbox"/> ("StoreID", "Hour")
<input type="checkbox"/> .range	<input type="checkbox"/> ("StoreID", "Year", "Month", "Day", "Hour")
<input type="checkbox"/> .sortBy	<input type="checkbox"/> ("Year", "Month", "Day", "Hour", "StoreID")

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

<input type="checkbox"/>
<input type="checkbox"/> .csv("/Purchases")
<input type="checkbox"/> .json("/Purchases")
<input type="checkbox"/> .parquet("/Purchases")
<input type="checkbox"/> .saveAsTable("/Purchases")

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 10

Topic #: 1

[\[All DP-300 Questions\]](#)

You are designing a streaming data solution that will ingest variable volumes of data.

You need to ensure that you can change the partition count after creation.

Which service should you use to ingest the data?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 11

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

You are building a database in an Azure Synapse Analytics serverless SQL pool.

You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container.

Records are structured as shown in the following sample.

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

The records contain two applicants at most.

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

CREATE EXTERNAL TABLE
CREATE TABLE
CREATE VIEW

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

CROSS APPLY
OPENJSON
OPENROWSET

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

GO

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 12

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure Synapse Analytics Apache Spark pool named Pool1.

You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file.

You need to load the files into the tables. The solution must maintain the source data types.

What should you do?

- A. Load the data by using PySpark.
- B. Load the data by using the OPENROWSET Transact-SQL command in an Azure Synapse Analytics serverless SQL pool.
- C. Use a Get Metadata activity in Azure Data Factory.
- D. Use a Conditional Split transformation in an Azure Synapse data flow.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 13

Topic #: 1

[\[All DP-300 Questions\]](#)

You are designing a date dimension table in an Azure Synapse Analytics dedicated SQL pool. The date dimension table will be used by all the fact tables. Which distribution type should you recommend to minimize data movement?

- A. HASH
- B. REPLICATE
- C. ROUND_ROBIN

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 14

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

From a website analytics system, you receive data extracts about user interactions such as downloads, link clicks, form submissions, and video plays.

The data contains the following columns:

Name	Sample value
Date	15 Jan 2021
EventCategory	Videos
EventAction	Play
EventLabel	Contoso Promotional
ChannelGrouping	Social
TotalEvents	150
UniqueEvents	120
SessionsWithEvents	99

You need to design a star schema to support analytical queries of the data. The star schema will contain four tables including a date dimension.

To which table should you add each column? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

EventCategory: ▼

DimChannel
DimDate
DimEvent
FactEvents

ChannelGrouping: ▼

DimChannel
DimDate
DimEvent
FactEvents

TotalEvents: ▼

DimChannel
DimDate
DimEvent
FactEvents

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 15

Topic #: 1

[\[All DP-300 Questions\]](#)

DRAG DROP -

You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.

Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted.

You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the amount of time required to delete old data.

How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not 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.

Select and Place:

Values

CustomerKey

HASH

ROUND_ROBIN

REPLICATE

OrderDateKey

SalesOrderNumber

Answer Area

```
CREATE TABLE [dbo].[FactSales]
(
    [ProductKey] int NOT NULL
, [OrderDateKey] int NOT NULL
, [CustomerKey] int NOT NULL
, [SalesOrderNumber] nvarchar ( 20 ) NOT NULL
, [OrderQuantity] smallint NOT NULL
, [UnitPrice] money NOT NULL
)
WITH
(
    CLUSTERED COLUMNSTORE INDEX
, DISTRIBUTION = [ ] ([ProductKey])
, PARTITION ( [ ] ) RANGE RIGHT FOR VALUES
(20170101, 20180101, 20190101, 20200101, 20210101)
)
```

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 16

Topic #: 1

[\[All DP-300 Questions\]](#)

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

You plan to create a database named DB1 in Pool1.

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

Which format should you use for the tables in DB1?

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 17

Topic #: 1

[\[All DP-300 Questions\]](#)

You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:

- ⇒ Send the output to an Azure Synapse.
- ⇒ Identify spikes and dips in time series data.
- ⇒ Minimize development and configuration effort.

Which should you include in the solution?

- A. Azure SQL Database
- B. Azure Databricks
- C. Azure Stream Analytics

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 18

Topic #: 1

[\[All DP-300 Questions\]](#)

You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL.

Which switch should you use to switch between languages?

A. \[<language>]

B. %<language>

C. \[<language>]

D. @<language>

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 19

Topic #: 1

[\[All DP-300 Questions\]](#)

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.



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 20

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

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

Git repository

Git repository information associated with your data factory. [CI/CD best practices](#)

Setting Disconnect

Repository type	Azure DevOps Git
Azure DevOps Account	CONTOSO
Project name	Data
Repository name	dwh_batchetl
Collaboration branch	main
Publish branch	adf_publish
Root folder	/

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.

Hot Area:

Answer Area

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

- /
- adf_publish
- main
- Parameterization template

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

- /contososales
- /dwh_batchetl/adf_publish/contososales
- /main

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 21

Topic #: 1

[\[All DP-300 Questions\]](#)

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

The output will be sent to a Delta Lake table.

Which output mode should you use?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 22

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.

You execute the Transact-SQL query shown in the following exhibit.

```
SELECT
    payment_type,
    SUM(fare_amount) AS fare_total
FROM OPENROWSET (
    BULK 'csv/busfare/tripdata_2020*.csv',
    DATA_SOURCE = 'BusData',
    FORMAT = 'CSV', PARSER_VERSION = '2.0',
    FIRSTROW = 2
)
WITH (
    payment_type INT 10,
    fare_amount FLOAT 11
) AS nyc
GROUP BY payment_type
ORDER BY payment_type;
```

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

Hot Area:

Answer Area

The query results include only **[answer choice]** in the csv/busfare folder.

	▼
CSV files in the tripdata_2020 subfolder	
files that have files names beginning with "tripdata_2020"	
CSV files that have file names containing "tripdata_202"	
CSV files that have file named beginning with "tripdata_2020"	

The query assumes that the first row in a CSV file is **[answer choice]** row.

	▼
a header	
a data	
an empty	

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 23

Topic #: 1

[\[All DP-300 Questions\]](#)

You have a SQL pool in Azure Synapse that contains a table named `dbo.Customers`. The table contains a column name `Email`.

You need to prevent nonadministrative users from seeing the full email addresses in the `Email` column. The users must see values in a format of `aXXX@XXXX.com` instead.

What should you do?

- A. From the Azure portal, set a mask on the `Email` column.
- B. From the Azure portal, set a sensitivity classification of `Confidential` for the `Email` column.
- C. From Microsoft SQL Server Management Studio, set an email mask on the `Email` column.
- D. From Microsoft SQL Server Management Studio, grant the `SELECT` permission to the users for all the columns in the `dbo.Customers` table except `Email`.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 24

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure Databricks workspace named workspace1 in the Standard pricing tier. Workspace1 contains an all-purpose cluster named cluster1.

You need to reduce the time it takes for cluster1 to start and scale up. The solution must minimize costs.

What should you do first?

- A. Upgrade workspace1 to the Premium pricing tier.
- B. Configure a global init script for workspace1.
- C. Create a pool in workspace1.
- D. Create a cluster policy in workspace1.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 25

Topic #: 1

[\[All DP-300 Questions\]](#)

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: In an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 26

Topic #: 1

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 27

Topic #: 1

[\[All DP-300 Questions\]](#)

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 a dedicated SQL pool to create an external table that has an additional DateTime column.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 28

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

Microsoft Azure Search resources, services, and docs (G+)

Home > SQL databases > Create SQL Database >

Configure

Feedback

Compute Hardware
Click "Change configuration" to see details for all hardware generations available including memory optimized and compute optimized options

Hardware Configuration

Gen5
up to 40 vCores, up to 120 GB memory
[Change configuration](#)

Max vCores: 1 2 4 6 8 10 12 14 16 18 20 24 32 40 6 vCores

Min vCores: 0.75 1 1.25 1.5 1.75 2 2.25 2.5 3 4 5 6 0.75 vCores

2.25 GB MIN MEMORY 18 GB MAX MEMORY

Auto-pause delay
The database automatically pauses if it is inactive for the time period specified here, and automatically resumes when database activity recurs. Alternatively, auto-pausing can be disabled.

Enable auto-pause

Days: 0 Hours: 4 Minutes: 0

Data max size: 1 GB 800 GB 1.5 TB 800 GB

240 GB LOG SPACE ALLOCATED

Apply

Cost summary

Gen5 - General Purpose (GP_5, Gen5_0)
Cost per GB (in USD) 0.12
Max storage selected (in GB) x 1040

ESTIMATED STORAGE COST / MONTH 119.60 USD
COMPUTE COST / VCORE / SECOND¹ 0.000145 USD

NOTES
¹ Serverless databases are billed in vCores based on a combination of CPU and memory utilization. [Learn more about serverless billing](#)

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.

Hot Area:

Answer Area

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

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 29

Topic #: 1

[\[All DP-300 Questions\]](#)

You plan to deploy an app that includes an Azure SQL database and an Azure web app. The app has the following requirements:

- ⇒ The web app must be hosted on an Azure virtual network.
- ⇒ The Azure SQL database must be assigned a private IP address.
- ⇒ The Azure SQL database must allow connections only from a specific virtual network.

You need to recommend a solution that meets the requirements.

What should you include in the recommendation?

- A. Azure Private Link
- B. a network security group (NSG)
- C. a database-level firewall
- D. a server-level firewall

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 30

Topic #: 1

[\[All DP-300 Questions\]](#)

You are planning a solution that will use Azure SQL Database. Usage of the solution will peak from October 1 to January 1 each year.

During peak usage, the database will require the following:

- ⇒ 24 cores
- ⇒ 500 GB of storage
- ⇒ 124 GB of memory
- ⇒ More than 50,000 IOPS

During periods of off-peak usage, the service tier of Azure SQL Database will be set to Standard.

Which service tier should you use during peak usage?

- A. Business Critical
- B. Premium
- C. Hyperscale

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 31

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have an Azure subscription.

You need to deploy an Azure SQL resource that will support cross database queries by using an Azure Resource Manager (ARM) template.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
"resources": [  
  ...  
  "type":   
    Microsoft.Sql/servers  
    Microsoft.Sql/servers/databases  
    Microsoft.Sql/managedInstances  
  "name": "[parameters('targetName')]",  
  "location": "[parameters('location')]",  
  "sku": {  
    "name": "[parameters('skuName')]"  
  }  
  ...  
  "dependsOn": [  
      
    "[parameters('virtualNetworkName')]"  
    "[variables('networkSecurityGroupName')]"  
  ],  
  "properties": {  
    "administratorLogin": "[parameters('administratorLogin')]",  
    "administratorLoginPassword": "[parameters('administratorLoginPassword')]",  
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets', parameters('virtualNetworkName'),  
    parameters('virtualNetworkName'), parameters('subnetName'))]",  
    "storageSizeInGB": "[parameters('storageSizeInGB')]", "vCores": "[parameters('vCores')]",  
    "licenseType": "[parameters('licenseType')]"  
  }  
  ...  
]
```

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 32

Topic #: 1

[\[All DP-300 Questions\]](#)

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"/>

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 33

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have an on-premises Microsoft SQL Server 2019 instance that hosts a database named DB1.

You plan to perform an online migration of DB1 to an Azure SQL managed instance by using the Azure Database Migration Service.

You need to create a backup of DB1 that is accessible to the Azure Database Migration Service.

What should you run for the backup and where should you store the backup? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Run:

A full backup and a log backup appended to the same file by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH FILE_SNAPSHOT option

Store the backup in:

A Recovery Services vault
An Azure Blob storage account
An SMB file share

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 34

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT -

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.

Hot Area:

Answer Area

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

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 35

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1.

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1.

The on-premises network connects to VNET1 by using an ExpressRoute connection.

You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service.

You need to configure VNET1 to support the migration.

What should you do?

- A. Configure service endpoints.
- B. Configure virtual network peering.
- C. Deploy an Azure firewall.
- D. Configure network security groups (NSGs).

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 36

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an on-premises Microsoft SQL server that uses the FileTables and Filestream features.

You plan to migrate to Azure SQL.

Which service should you use?

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 37

Topic #: 1

[\[All DP-300 Questions\]](#)

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

What should you do?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 38

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure SQL database named DB1.

You have a table name Table1 that has 20 columns of type CHAR(400). Row compression for Table1 is enabled.

During a database audit, you discover that none of the fields contain more than 150 characters.

You need to ensure that you can apply page compression to Table1.

What should you do?

- A. Configure the columns as sparse.
- B. Change the column type to NVARCHAR(MAX).
- C. Change the column type to VARCHAR(MAX).
- D. Change the column type to VARCHAR(200).

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 39

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an on-premises Microsoft SQL Server named SQL1 that hosts five databases.

You need to migrate the databases to an Azure SQL managed instance. The solution must minimize downtime and prevent data loss.

What should you use?

- A. Always On availability groups
- B. Backup and Restore
- C. log shipping
- D. Database Migration Assistant

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 40

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure subscription that contains an Azure SQL database. The database contains a table named `table1` that uses partitioned columnstores.

You need to configure `table1` to meet the following requirements:

- Each partition must be compressed.
- The compression ratio must be maximized.
- You must be able to index the compressed data.

What should you use?

- A. page compression
- B. columnstore compression
- C. GZIP compression
- D. columnstore archival compression

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 41

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure subscription linked to an Azure Active Directory (Azure AD) tenant. The subscription contains 10 virtual machines that run Windows Server 2019 and host Microsoft SQL Server 2019 instances.

You need to ensure that you can manage the SQL Server instances by using a single user account.

What should you do first?

- A. Enable a user-assigned managed identity on each virtual machine.
- B. Deploy an Azure Active Directory Domain Services (Azure AD DS) domain and join the virtual machines to the domain.
- C. Enable a system-assigned managed identity on each virtual machine.
- D. Join the virtual machines to the Azure AD tenant.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 42

Topic #: 1

[\[All DP-300 Questions\]](#)

DRAG DROP

-

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 throughput 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, more than once, or not 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:

Transaction logs:

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 43

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure SQL Database elastic pool that contains 10 databases.

You receive the following alert.

Msg 1132, Level 16, state 1, Line 1

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

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

Which three actions can you perform? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 44

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure subscription.

You need to deploy a new Azure SQL database by using Azure Command-Line Interface (CLI).

Which three parameters are required?

- A. --name, --edition, and --capacity
- B. --name, --tier, and --min-capacity
- C. --name, --resource-group, and --server
- D. --name, --licence-type, and --capacity

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 45

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT

-

You have an Azure subscription.

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

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

The solution must minimize administrative effort.

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

NOTE: Each correct selection is worth one point.

Answer Area

MSX:

	▼
SQL database	
SQL managed instances	
SQL virtual machines	

TSX:

	▼
SQL database	
SQL managed instances	
SQL virtual machines	

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 46

Topic #: 1

[\[All DP-300 Questions\]](#)

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

You need to migrate the databases hosted on SQL1 to Azure. The solution must meet the following requirements:

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

What should you use to host the databases?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 47

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT

You have an on-premises Microsoft SQL Server 2016 instance that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named MI1.

You plan to perform an online migration of db1 to MI1 by using Azure Database Migration Service.

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

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

NOTE: Each correct selection is worth one point.

Answer Area

Backup type:

	▼
Full only	
Full and differential	
Full and transaction log	
Transaction log only	

To store the backups:

	▼
Append all the backups to a single file.	
Append each backup to a separate file.	
Write each backup to a separate file.	

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 48

Topic #: 1

[\[All DP-300 Questions\]](#)

You have a SQL Server on Azure Virtual Machines instance named SQLVM1 that was deployed by using an Azure Marketplace SQL Server 2019 Enterprise image.

You need to change the Microsoft SQL Server instance on SQLVM1 to the Standard edition. The solution must ensure licensing compliance.

What should you do first?

- A. From the SQL Server Installation Center on SQLVM1, run the Edition Upgrade wizard.
- B. From SQLVM1, uninstall the SQL Server instance.
- C. From the SQL Server Installation Center on SQLVM1, run the Repair wizard.
- D. From the Azure portal, reconfigure SQLVM1.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 49

Topic #: 1

[\[All DP-300 Questions\]](#)

Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database named db1.

You have an Azure subscription.

You plan to migrate db1 to an Azure SQL managed instance.

You need to create the SQL managed instance. The solution must minimize the disk latency of the instance.

Which service tier should you use?

- A. Business Critical
- B. Hyperscale
- C. General Purpose
- D. Premium

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 50

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure subscription.

You need to deploy an Azure SQL database. The solution must meet the following requirements:

- Dynamically scale CPU resources.
- Ensure that the database can be paused to reduce costs.

What should you use?

- A. the Business Critical service tier
- B. the serverless compute tier
- C. an elastic pool
- D. the General Purpose service tier

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 51

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT

-

You have an Azure subscription.

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

Service tier:

	▼
Business Critical	
General Purpose	
Hyperscale	

Hardware generation:

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 52

Topic #: 1

[\[All DP-300 Questions\]](#)

You have a Microsoft SQL Server 2017 server.

You need to migrate the server to Azure. The solution must meet the following requirements:

- Ensure that the latest version of SQL Server is used.
- Support the SQL Server Agent service.
- Minimize administrative effort.

What should you use?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 53

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT

-

You have a Microsoft SQL Server 2017 server that hosts five databases.

You plan to migrate the databases to Azure.

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

Azure service:

- A single Azure SQL database in the provisioned compute tier
- A single Azure SQL database in the serverless computer tier
- An Azure SQL Database elastic pool
- Azure SQL Managed Instance

Service tier:

- Basic
- General Purpose
- Standard

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 54

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an on-premises Microsoft SQL Server 2019 database named SQL1 that uses merge replication.

You need to migrate SQL1 to Azure.

Which service should you use?

- A. Azure SQL Edge
- B. Azure SQL Database
- C. SQL Server on Azure Virtual Machines
- D. Azure SQL Managed Instance

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 55

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an on-premises datacenter that contains a 2-TB Microsoft SQL Server 2019 database named DB1.

You need to recommend a solution to migrate DB1 to an Azure SQL managed instance. The solution must minimize downtime and administrative effort.

What should you include in the recommendation?

- A. Log Replay Service (LRS)
- B. log shipping
- C. transactional replication
- D. SQL Data Sync

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 56

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure subscription.

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

Which virtual machine series should you use?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 57

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an on-premises Microsoft SQL Server 2019 instance that hosts a database named DB1.

You have an Azure subscription that contains an Azure SQL database named SQLDB1.

You need to replicate DB1 to SQLDB1.

Which type of replication should you use?

- A. transactional
- B. peer-to-peer
- C. snapshot
- D. merge

[Show Suggested Answer](#)



Actual exam question from Microsoft's DP-300

Question #: 58

Topic #: 1

[\[All DP-300 Questions\]](#)

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.

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

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 59

Topic #: 1

[\[All DP-300 Questions\]](#)

DRAG DROP

-

You have an Azure SQL database named DB1.

You need to create a partitioned table in DB1.

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.

Objects

a partition scheme

an aligned index

a filegroup

a table

a partition function

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 60

Topic #: 1

[\[All DP-300 Questions\]](#)

DRAG DROP

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a virtual machine named Server1 that runs Windows Server. SQLVM1 and Server1 are joined to an Active Directory Domain Services (AD DS) domain. Server1 hosts a file share named Share1.

You need to ensure that a SQL Server Agent job step on SQLVM1 can access the files in Share1. The solution must use the principle of least privilege.

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

Assign the proxy to the job step.

Create a login.

Create a database user.

Create a credential.

Create a proxy.



Answer Area



Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 61

Topic #: 1

[\[All DP-300 Questions\]](#)

You have an Azure subscription.

You need to deploy an instance of SQL Server on Azure Virtual Machines. The solution must meet the following requirements:

- Custom performance configuration, such as IOPS, capacity, and throughout, must be supported.
- Costs must be minimized.

Which type of disk should you include in the solution?

- A. Premium SSD v2
- B. Premium SSD
- C. Standard SSD
- D. Ultra SSD

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 62

Topic #: 1

[\[All DP-300 Questions\]](#)

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

You plan to create an Azure SQL managed instance and migrate the on-premises database to the new instance.

Which three service tiers support the SQL managed instance? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 68

Topic #: 1

[\[All DP-300 Questions\]](#)

SIMULATION

-

You need to add an Azure AD user named user2-12345678@examusers.com to db1. User2-12345678 must be able to read data from all the tables in db1 without being able to modify the data.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

[Show Suggested Answer](#)



Actual exam question from Microsoft's DP-300

Question #: 70

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT

-

Case study

-

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

-

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

-

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment. Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Existing Environment. Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Existing Environment. Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Existing Environment. Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements. Planned Changes

-

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Requirements. Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Requirements. Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Requirements. Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Requirements. Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to recommend which service and target endpoint to use when migrating the databases from SVR1 to Instance1. The solution must meet the availability requirements.

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

NOTE: Each correct selection is worth one point.

Answer Area

Migration service:

Target endpoint:

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 71

Topic #: 1

[\[All DP-300 Questions\]](#)

HOTSPOT

-

Case study

-

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

-

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

-

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment. Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Existing Environment. Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Existing Environment. Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Existing Environment. Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements. Planned Changes

-

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Requirements. Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Requirements. Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Requirements. Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Requirements. Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to recommend a service tier and a method to offload analytical workloads for the databases migrated from SVR1. The solution must meet the availability and business requirements.

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

NOTE: Each correct selection is worth one point.

Answer Area

Service tier:

Method:

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 1

Topic #: 2

[\[All DP-300 Questions\]](#)

You have a new Azure SQL database. The database contains a column that stores confidential information.

You need to track each time values from the column are returned in a query. The tracking information must be stored for 365 days from the date the query was executed.

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

NOTE: Each correct selection is worth one point.

- A. Turn on auditing and write audit logs to an Azure Storage account.
- B. Add extended properties to the column.
- C. Turn on auditing and write audit logs to an Event Hub
- D. Apply sensitivity labels named Highly Confidential to the column.
- E. Turn on Azure Defender for SQL

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 2

Topic #: 2

[\[All DP-300 Questions\]](#)

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

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

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

- ⇒ Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.
- ⇒ Minimize the possibility of VM1 exfiltrating data stored in SqlDb1.

What should you create on VNet1?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 3

Topic #: 2

[\[All DP-300 Questions\]](#)

DRAG DROP -

You have a new Azure SQL database named DB1 on an Azure SQL server named AzSQL1.

The only user who was created is the server administrator.

You need to create a contained database user in DB1 who will use Azure Active Directory (Azure AD) for authentication.

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

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Connect to DB1 by using the server administrator account.

Set the Active Directory Admin for AzSQL1.

From the Azure portal, assign the SQL DB Contributor role to the user.

Create a login in the master database.



Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 4

Topic #: 2

[\[All DP-300 Questions\]](#)

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:

- ⇒ 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.

Hot Area:

Answer Area

Exposed Prefix:

Padding String:

Exposed Suffix:

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 5

Topic #: 2

[\[All DP-300 Questions\]](#)

DRAG DROP -

You have an Azure SQL database that contains a table named Employees. Employees contains a column named Salary.

You need to encrypt the Salary column. The solution must prevent database administrators from reading the data in the Salary column and must provide the most secure encryption.

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

Select and Place:

Actions

Encrypt the Salary column by using the randomized encryption type.

Create a column encryption key.

Enable Transparent Data Encryption (TDE).

Encrypt the Salary column by using the deterministic encryption type.

Apply a dynamic data mask to the Salary column.

Create a column master key.

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 6

Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have an Azure SQL database named DB1 that contains two tables named Table1 and Table2. Both tables contain a column named a Column1. Column1 is used for joins by an application named App1.

You need to protect the contents of Column1 at rest, in transit, and in use.

How should you protect the contents of Column1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Encryption key:

	▼
Column encryption key	
Database encryption key	
Service master key	

Encryption type:

	▼
Deterministic	
Randomized	
Transparent Data Encryption (TDE)	

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 7

Topic #: 2

[\[All DP-300 Questions\]](#)

You have 40 Azure SQL databases, each for a different customer. All the databases reside on the same Azure SQL Database server.

You need to ensure that each customer can only connect to and access their respective database.

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

NOTE: Each correct selection is worth one point.

- A. Implement row-level security (RLS).
- B. Create users in each database.
- C. Configure the database firewall.
- D. Configure the server firewall.
- E. Create logins in the master database.
- F. Implement Always Encrypted.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 8

Topic #: 2

[\[All DP-300 Questions\]](#)

DRAG DROP -

You have an Azure SQL Database instance named DatabaseA on a server named Server1.

You plan to add a new user named App1 to DatabaseA and grant App1 db_datareader permissions. App1 will use SQL Server Authentication.

You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.

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

On the master database, run `CREATE LOGIN [APP1] FROM EXTERNAL PROVIDER;`

On DatabaseA, run `CREATE USER [APP1] WITH PASSWORD = 'P@ssW0rd!';`

On DatabaseA, run `ALTER ROLE db_datareader ADD MEMBER [App1];`

On the master database, run `CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';`

On DatabaseA, run `CREATE USER [App1] FROM LOGIN [App1];`

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 9

Topic #: 2

[\[All DP-300 Questions\]](#)

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

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

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

- ⇒ Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1.
- ⇒ Restrict network connectivity to SqlSrv1.

What should you create on VNet1?

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 10

Topic #: 2

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 11

Topic #: 2

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 12

Topic #: 2

[\[All DP-300 Questions\]](#)

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	

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 13

Topic #: 2

[\[All DP-300 Questions\]](#)

You have a data warehouse in Azure Synapse Analytics.

You need to ensure that the data in the data warehouse is encrypted at rest.

What should you enable?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 14

Topic #: 2

[\[All DP-300 Questions\]](#)

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

You need to ensure that users from each company can view only the data of their respective company.

Which two objects should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 15

Topic #: 2

[\[All DP-300 Questions\]](#)

You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1. DF1 contains a linked service.

You have an Azure Key vault named vault1 that contains an encryption key named key1.

You need to encrypt df1 by using key1.

What should you do first?

- A. Disable purge protection on vault1.
- B. Remove the linked service from df1.
- C. Create a self-hosted integration runtime.
- D. Disable soft delete on vault1.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 16

Topic #: 2

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 17

Topic #: 2

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 18

Topic #: 2

[\[All DP-300 Questions\]](#)

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

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

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

- ⇒ Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1.
- ⇒ 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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 19

Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have an Azure SQL database named db1 that contains an Azure Active Directory (Azure AD) user named user1.

You need to test impersonation of user1 in db1 by running a SELECT statement and returning to the original execution context.

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.

Hot Area:

Answer Area

EXECUTE AS

	▼
CALLER	
LOGIN	
OWNER	
USER	

 = 'user1@contoso.com'

GO

SELECT SUSER_SNAME ()

	▼
REVERT	
REVOKE	
ROLLBACK	

GO

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 20

Topic #: 2

[\[All DP-300 Questions\]](#)

DRAG DROP -

You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1.

You need to encrypt the data in Col1.

Which four actions should you perform for DB1 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

Create a database master key.

Create a column master key.

Open the symmetric key.

Create a certificate.

Update Col1.

Create a symmetric key.



Answer Area



Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 21

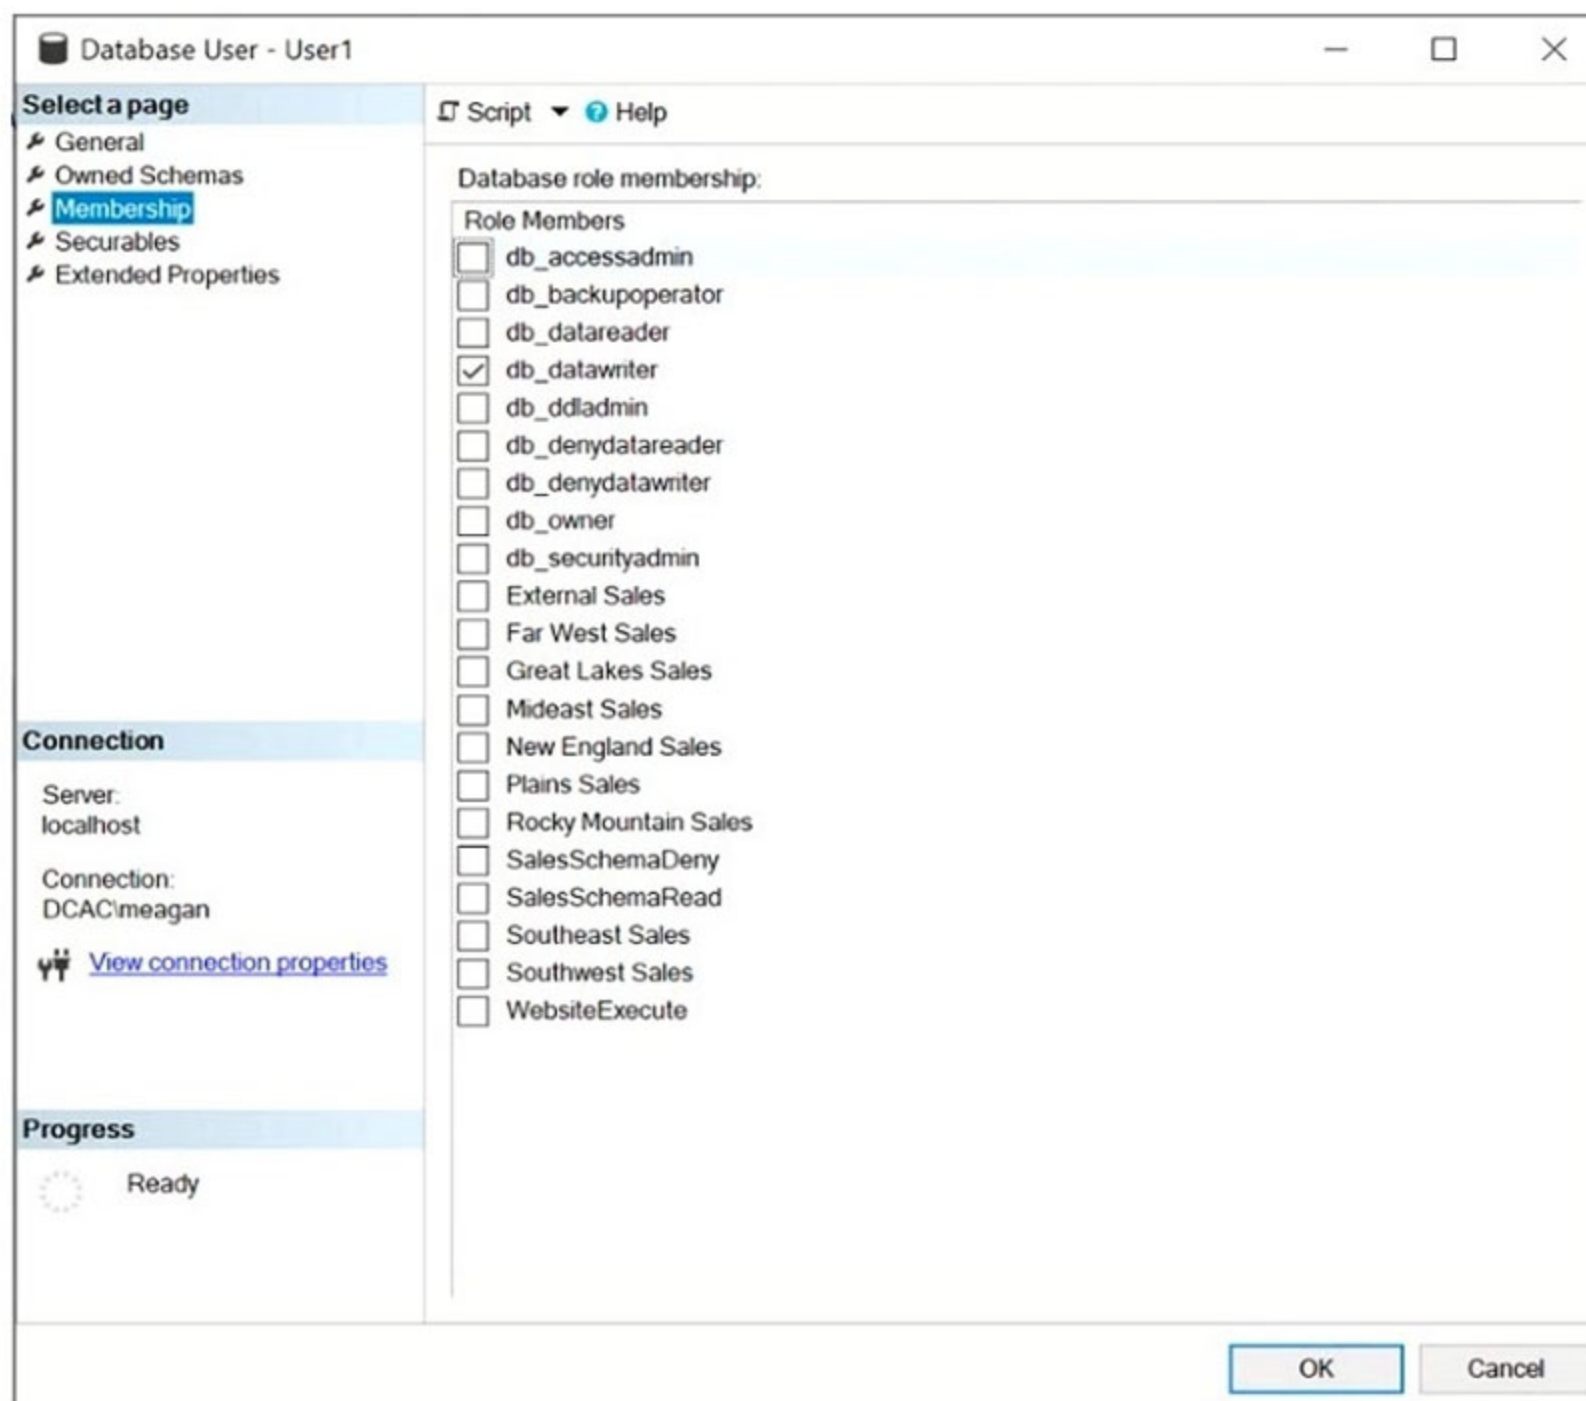
Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have a Microsoft SQL Server database named DB1 that contains a table named Table1.

The database role membership for a user named User1 is 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.

Hot Area:

Answer Area

User1 can **[answer choice]**.

	▼
add a column to Table1	
delete a row from Table1	
drop Table1	

To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the **[answer choice]** database role.

	▼
db_datareader	
db_ddladmin	
db_denydatareader	
db_denydatawriter	

Actual exam question from Microsoft's DP-300

Question #: 22

Topic #: 2

[\[All DP-300 Questions\]](#)

You have an Azure subscription that contains a logical SQL server named Server1. The master database of Server1 contains a user named User1.

You need to ensure that User1 can create databases on Server1.

Which database role should you assign to User1?

- A. db_owner
- B. dbmanager
- C. dbo
- D. db_ddladmin

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 23

Topic #: 2

[\[All DP-300 Questions\]](#)

You have an on-premises Microsoft SQL Server 2019 instance named SQL1 that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named MI1 and an Azure Storage account named storage1.

You plan to migrate db1 to MI1 by using the backup and restore process.

You need to ensure that you can back up db1 to storage1. The solution must meet the following requirements:

- ⇒ Use block blob storage.
- ⇒ Maximize security.

What should you do on storage1?

- A. Generate a shared access signature (SAS).
- B. Create an access policy.
- C. Rotate the storage keys.
- D. Enable infrastructure encryption.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 24

Topic #: 2

[\[All DP-300 Questions\]](#)

You have an Azure SQL database named DB1.

A user named User1 has an Azure Active Directory (Azure AD) account.

You need to provide User1 with the ability to add and remove columns from the tables in DB1. The solution must use the principle of least privilege.

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

NOTE: Each correct selection is worth one point.

- A. Assign the database user the db_owner role.
- B. Create a contained database user.
- C. Create a login and an associated database user.
- D. Assign the database user the db_ddladmin role.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 25

Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT

-

You have an Azure SQL database named DB1 that contains a table named Orders. The Orders table contains a row for each sales order. Each sales order includes the name of the user who placed the order.

You need to implement row-level security (RLS). The solution must ensure that the users can view only their respective sales orders.

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.

Create:

- A materialized view in DB1
- A security policy in the Orders table
- Database scoped credentials in DB1

Control access to the rows by using:

- A masking rule
- A table-valued function
- The CONTAINS predicate

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 26

Topic #: 2

[\[All DP-300 Questions\]](#)

DRAG DROP

You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1.

You plan to deploy an Azure web app named webapp1 that will export the rows in Table1 that have changed.

You need to ensure that webapp1 can identify the changes to Table1. The solution must meet the following requirements:

- Minimize compute times.
- Minimize storage.

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

Actions

From webapp1, connect to SQLDb1, obtain the initial dataset and run the `CHANGE_TRACKING_CURREWT_VERSION()` function.

Connect to SQLDb1 and run the following Transact-SQL statement
`EXEC sys.sp_cdc_enable_table`

Connect to SQLDb1 and run the following Transact-SQL statement
`EXEC sys.sp_cdc_enable_db`

From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGETABLE()` function.

Connect to SQLDb1 and run the following Transact-SQL statement
`ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON`

Connect to SQLDb1 and run the following Transact-SQL statement
`ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING`

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 27

Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT

You have an Azure subscription that is linked to an Azure Active Directory (Azure AD) tenant named contoso.com. The subscription contains an Azure SQL database named SQL1 and an Azure web app named app1. App1 has the managed identity feature enabled.

You need to create a new database user for app1.

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

NOTE: Each correct selection is worth one point.

Answer Area

CREATE USER FROM

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 28

Topic #: 2

[\[All DP-300 Questions\]](#)

DRAG DROP

You have an Azure subscription that contains an Azure SQL database named DB1.

You plan to perform a classification scan of DB1 by using Azure Purview.

You need to ensure that you can register DB1.

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

From the Purview Studio, create a scan trigger.

Modify the Access control (IAM) settings for DB1.

Create a system-assigned managed identity.

Create a Purview account.

From the Purview Studio, create a scan rule set.



Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 29

Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT

-

You have an Azure subscription that contains a group named Group1 and an Azure SQL managed instance that hosts a database named DB1.

You need to ensure that Group1 has read access to new tables created in DB1. The solution must use the principle of least privilege.

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

NOTE: Each correct selection is worth one point.

Answer Area

GRANT

	▼
DELETE	
INSERT	
SELECT	
UPDATE	

	▼
Database	
Schema	
Table	

:::[table] TO [Contoso\group1]

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 30

Topic #: 2

[\[All DP-300 Questions\]](#)

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

You plan to schedule a SQL Server Agent job that will rebuild indexes of the databases hosted on VM1.

You need to configure the account that will be used by the agent. The solution must use the principle of least privilege.

Which operating system user right should you assign to the account?

- A. Increase scheduling priority
- B. Log on as a service
- C. Profile system performance
- D. Log on as a batch job

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 31

Topic #: 2

[\[All DP-300 Questions\]](#)

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

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

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

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

[Show Suggested Answer](#)



Actual exam question from Microsoft's DP-300

Question #: 32

Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT

You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD service principal named app1.

You need to ensure that app1 can access db1. The solution must use the principle of least privilege.

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

NOTE: Each correct selection is worth one point.

Answer Area

CREATE

[app1]

- CREDENTIAL
- LOGIN
- USER

[app1]

- FOR LOGIN app1
- FROM EXTERNAL PROVIDER
- FROM LOGIN app1
- WITHOUT LOGIN

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 33

Topic #: 2

[\[All DP-300 Questions\]](#)

Your on-premises network contains a server that hosts a 60-TB database named DB1. The network has a 10-Mbps internet connection.

You need to migrate DB1 to Azure. The solution must minimize how long it takes to migrate the database.

What should you use?

- A. Azure Migrate
- B. Azure Data Box
- C. Azure Database Migration Service
- D. Data Migration Assistant (DMA)

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 34

Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have an Azure SQL logical server.

You run the following script.

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

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements

The `orders` table will allow only rows to be inserted.

Yes

No

To create additional tables in the Sales database, the `LEDGER = ON` parameter must be used.

To ensure that a timestamp is added to each row in the `orders` table, the `GENERATED ALWAYS` column must be added to the script.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 35

Topic #: 2

[\[All DP-300 Questions\]](#)

You have an Azure SQL database named DB1.

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

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

What should you include in the solution?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 36

Topic #: 2

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 37

Topic #: 2

[\[All DP-300 Questions\]](#)

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

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

You plan to enable Always Encrypted for Table1.

Which two columns support encryption? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 38

Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT

-

You have an Azure subscription that contains an Azure SQL managed instance named SQL1.

You use the Microsoft Power BI service.

You need to configure connectivity from the Power BI service to SQL1. The solution must ensure that only the Power BI service can initiate connections to SQL1.

Which type of endpoint should you use for SQL1, and what should you use to control the connections to SQL1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Endpoint type:

	▼
Private	
Public	
Service	

Control connections by using:

	▼
A database-level firewall	
A network security group (NSG)	
A server-level firewall	
Virtual network firewall rules	

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 39

Topic #: 2

[\[All DP-300 Questions\]](#)

HOTSPOT

-

You have an Azure subscription.

You need to deploy a logical SQL server by using PowerShell. The solution must ensure that the logical SQL server can create Azure AD users and provide Transparent Data Encryption (TDE) with a customer-managed key.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
New-AzSqlServer -ResourceGroupName "RG1" -ServerName "SQL1" -Location "EastUS" -ErrorAction Stop
```

```
-Tags @{Environment="Databases";Department="Data Tech"}
```

- assignidentity
- federatedclientID
- keyid

- assignidentity
- federatedclientID
- keyid

```
"https://db1.vault.azure.net/keys/dbkey/01234234512345678901234561823942"
```

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 40

Topic #: 2

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 41

Topic #: 2

[\[All DP-300 Questions\]](#)

DRAG DROP

You have an Azure key vault named Vault1 and a SQL server on Azure Virtual Machines instance named SQL1. SQL1 hosts a database named DB1.

You need to configure Transparent Data Encryption (TDE) on DB1 to use a key in Vault1.

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

On SQL1, create a cryptographic provider and a Microsoft SQL Server credential.

On SQL1, create a symmetric key.

On SQL1, create a login from the asymmetric key.

On SQL1, create a login from the symmetric key.

Create an Azure AD service principal and grant the service principal permissions for Vault1.

On SQL1, create an asymmetric key.

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 42

Topic #: 2

[\[All DP-300 Questions\]](#)

You have an Azure subscription that contains an instance of SQL Server on an Azure virtual machine named VM1 and an Azure Active Directory Domain Services (Azure AD DS) domain that contains two users named User1 and User 2.

On the default instance of SQL Server on VM1, you create a credential named Credential1 for User1.

You need to ensure that User2 can create a SQL Server Agent proxy that will use Credential1. The solution must use the principle of least privilege.

Which role should you assign to User2?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 43

Topic #: 2

[\[All DP-300 Questions\]](#)

DRAG DROP

You have an instance of SQL Server on Azure Virtual Machines named SQL1. SQL1 contains a database named DB1.

You need to enable Transparent Data Encryption (TDE) for DB1.

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.

Objects

a database encryption key in the master database

a master key in the master database

a certificate in DB1

a certificate in the master database

a master key in DB1

a database encryption key in DB1

Answer Area

1

2

3



Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 45

Topic #: 2

[\[All DP-300 Questions\]](#)

SIMULATION

-

You need to configure your user account as the Azure AD admin for the server named sql12345678.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

[Show Suggested Answer](#)





Actual exam question from Microsoft's DP-300

Question #: 46

Topic #: 2

[\[All DP-300 Questions\]](#)

SIMULATION

-

You need to create a new database role named role1 that can read the contents of the SalesLT.SalesOrderDetail table only in db1.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

[Show Suggested Answer](#)





Actual exam question from Microsoft's DP-300

Question #: 47

Topic #: 2

[\[All DP-300 Questions\]](#)

SIMULATION

-

You need to classify the StandardCost and ListPrice columns of the SalesLT.Product table in db1. The columns must have the following configurations:

- An information type of Financial
- A sensitivity label of Confidential

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

[Show Suggested Answer](#)





Actual exam question from Microsoft's DP-300

Question #: 48

Topic #: 2

[\[All DP-300 Questions\]](#)

SIMULATION

-

You need to protect all the databases on sql12345678 from SQL injection attacks.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

[Show Suggested Answer](#)



Actual exam question from Microsoft's DP-300

Question #: 49

Topic #: 2

[\[All DP-300 Questions\]](#)

Case study -

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview -

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment. Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Existing Environment. Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Existing Environment. Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Existing Environment. Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements. Planned Changes -

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Requirements. Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Requirements. Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Requirements. Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Requirements. Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to recommend a process to automate the management of DB3. The solution must meet the management requirements.

What should be the first step of the process?

- Configure Microsoft Entra authentication for the logical server that hosts DB3.
- Configure a private endpoint for connectivity to DB3.
- Create database-scoped credentials in DB3.
- Create a database that has database-scoped credentials.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 51

Topic #: 2

[\[All DP-300 Questions\]](#)

Case study -

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview -

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment. Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Existing Environment. Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Existing Environment. Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Existing Environment. Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements. Planned Changes -

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Requirements. Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Requirements. Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Requirements. Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Requirements. Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to recommend a solution that will enable remote developers to access DB1 and DB2. The solution must support the planned changes and meet the security requirements.

What should you include in the recommendation?

- A. a public endpoint via a database-level firewall rule
- B. a Point-to-Site (P2S) VPN
- C. a public endpoint via a server-level firewall rule
- D. a private endpoint

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 1

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure SQL database named sqldb1.

You need to minimize the possibility of Query Store transitioning to a read-only state.

What should you do?

- A. Double the value of Data Flush interval
- B. Decrease by half the value of Data Flush Interval
- C. Double the value of Statistics Collection Interval
- D. Decrease by half the value of Statistics Collection interval

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 2

Topic #: 3

[\[All DP-300 Questions\]](#)

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 16 vCPUs 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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 3

Topic #: 3

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have an Azure SQL database named db1.

You need to retrieve the resource usage of db1 from the last week.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

SELECT *

FROM

	▼
sys.dm_db_resource_stats	
sys.dm_exec_requests	
sys.dm_user_db_resource_governance	
sys.resource_stats	

WHERE database_name = 'db1' AND

start_time >

	▼
DATEADD	
DATEDIFF	
DATEPART	
TODATETIMEOFFSET	

(day, -7, GETDATE())

ORDER BY start_time DESC;

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 4

Topic #: 3

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 5

Topic #: 3

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 6

Topic #: 3

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 7

Topic #: 3

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 8

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure virtual machine based on a custom image named VM1.

VM1 hosts an instance of Microsoft SQL Server 2019 Standard.

You need to automate the maintenance of VM1 to meet the following requirements:

- ⇒ Automate the patching of SQL Server and Windows Server.
- ⇒ Automate full database backups and transaction log backups of the databases on VM1.
- ⇒ Minimize administrative effort.

What should you do first?

- A. Enable a system-assigned managed identity for VM1
- B. Register the Azure subscription to the Microsoft.Sql resource provider
- C. Install an Azure virtual machine Desired State Configuration (DSC) extension on VM1
- D. Register the Azure subscription to the Microsoft.SqlVirtualMachine resource provider

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 9

Topic #: 3

[\[All DP-300 Questions\]](#)

HOTSPOT -

You are building an Azure Stream Analytics job to retrieve game data.

You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game.

How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

SELECT as HighestScore

- Collect(Score)
- CollectTop(1)OVER(ORDER BY Score Desc)
- Game, MAX(Score)
- TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)

FROM input TIMESTAMP BY CreatedAt

GROUP BY

- Game
- Hopping(minute, 5)
- Tumbling(minute, 5)
- Windows(TumblingWindow(minute, 5), Hopping(minute, 5))

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 10

Topic #: 3

[\[All DP-300 Questions\]](#)

A company plans to use Apache Spark analytics to analyze intrusion detection data.

You need to recommend a solution to analyze network and system activity data for malicious activities and policy violations. The solution must minimize administrative efforts.

What should you recommend?

- A. Azure Data Lake Storage
- B. Azure Databricks
- C. Azure HDInsight
- D. Azure Data Factory

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 11

Topic #: 3

[\[All DP-300 Questions\]](#)

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



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 12

Topic #: 3

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 13

Topic #: 3

[\[All DP-300 Questions\]](#)

You are monitoring an Azure Stream Analytics job.

You discover that the Backlogged input Events metric is increasing slowly and is consistently non-zero.

You need to ensure that the job can handle all the events.

What should you do?

- A. Remove any named consumer groups from the connection and use \$default.
- B. Change the compatibility level of the Stream Analytics job.
- C. Create an additional output stream for the existing input stream.
- D. Increase the number of streaming units (SUs).

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 14

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure Stream Analytics job.

You need to ensure that the job has enough streaming units provisioned.

You configure monitoring of the SU % Utilization metric.

Which two additional metrics should you monitor? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 15

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure Databricks resource.

You need to log actions that relate to changes in compute for the Databricks resource.

Which Databricks services should you log?

- A. clusters
- B. jobs
- C. DBFS
- D. SSH
- E. workspace

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 16

Topic #: 3

[\[All DP-300 Questions\]](#)

Your company uses Azure Stream Analytics to monitor devices.

The company plans to double the number of devices that are monitored.

You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load.

Which metric should you monitor?

- A. Input Deserialization Errors
- B. Late Input Events
- C. Early Input Events
- D. Watermark delay

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 17

Topic #: 3

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 18

Topic #: 3

[\[All DP-300 Questions\]](#)

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

You need to identify the extent of the data skew in Table1.

What should you do in Synapse Studio?

- A. Connect to Pool1 and query `sys.dm_pdw_nodes_db_partition_stats`.
- B. Connect to the built-in pool and run `DBCC CHECKALLOC`.
- C. Connect to Pool1 and run `DBCC CHECKALLOC`.
- D. Connect to the built-in pool and query `sys.dm_pdw_nodes_db_partition_stats`.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 19

Topic #: 3

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 20

Topic #: 3

[\[All DP-300 Questions\]](#)

DRAG DROP -

You have an Azure SQL managed instance named SQLMI1 that has Resource Governor enabled and is used by two apps named App1 and App2.

You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1.

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

Create a workload group.

Create a user-defined classifier function.

Modify Resource Governor.

Create a contained database user.

Create a resource pool.

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 21

Topic #: 3

[\[All DP-300 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance.

Solution: You shrink the transaction log file.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 22

Topic #: 3

[\[All DP-300 Questions\]](#)

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 change the data file for the master database to autogrow by 10 percent.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 23

Topic #: 3

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 24

Topic #: 3

[\[All DP-300 Questions\]](#)

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

You need to modify the MAXDOP settings for db1.

What should you do?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 25

Topic #: 3

[\[All DP-300 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance.

Solution: You create additional tempdb files.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 26

Topic #: 3

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 27

Topic #: 3

[\[All DP-300 Questions\]](#)

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

The Intelligent Insights diagnostics log identifies that several tables are missing indexes.

You need to ensure that indexes are created for the tables.

What should you do?

- A. Run the DBCC SQLPERF command.
- B. Run the DBCC DBREINDEX command.
- C. Modify the automatic tuning settings for db1.
- D. Modify the Query Store settings for db1.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 28

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2.

You need to limit the number of IOPs that App2 queries generate on SQL1.

Which two actions should you perform on SQL1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable query optimizer fixes.
- B. Enable Resource Governor.
- C. Enable parameter sniffing.
- D. Create a workload group.
- E. Configure In-memory OLTP.
- F. Run the Database Engine Tuning Advisor.
- G. Reduce the Max Degree of Parallelism value.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 29

Topic #: 3

[\[All DP-300 Questions\]](#)

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

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

You need to resolve the performance issues.

What should you do?

- A. Implement memory-optimized tables.
- B. Run the DBCC FLUSHPROCINDB command.
- C. Replace the sequential index keys with nonsequential keys.
- D. Run the DBCC DBREINDEX command.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 30

Topic #: 3

[\[All DP-300 Questions\]](#)

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 dmvs2

ON dmvs1.session_id = dmvs2.session_id

WHERE is_user_process = 1

GROUP BY wait_type

ORDER BY SUM(wait_time) DESC;

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 31

Topic #: 3

[\[All DP-300 Questions\]](#)

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 vCPUs and 64 GB of memory.

You need to reduce tempdb contention without negatively affecting server performance.

What is the number of secondary data files that you should configure for tempdb?

- A. 2
- B. 4
- C. 8
- D. 64

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 32

Topic #: 3

[\[All DP-300 Questions\]](#)

You receive numerous alerts from Azure Monitor for an Azure SQL Database instance.

You need to reduce the number of alerts. You must only receive alerts if there is a significant change in usage patterns for an extended period.

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

NOTE: Each correct selection is worth one point.

- A. Set Threshold Sensitivity to High
- B. Set the Alert logic threshold to Dynamic
- C. Set the Alert logic threshold to Static
- D. Set Threshold Sensitivity to Low
- E. Set Force Plan to On

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 33

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure SQL database named sqlldb1.

You need to minimize the amount of space by the data and log files of sqlldb1.

What should you run?

- A. DBCC SHRINKDATABASE
- B. sp_clean_db_free_space
- C. sp_clean_db_file_free_space
- D. DBCC SHRINKFILE

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 34

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure SQL Database server named sqlsrv1 that hosts 10 Azure SQL databases.

The databases perform slower than expected.

You need to identify whether the performance issue relates to the use of tempdb by Azure SQL databases in sqlsrv1.

What should you do?

- A. Run Query Store-based queries
- B. Review information provided by SQL Server Profiler-based traces
- C. Review information provided by Query Performance Insight
- D. Run dynamic management view-based queries

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 35

Topic #: 3

[\[All DP-300 Questions\]](#)

DRAG DROP -

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.

Select and Place:

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

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 36

Topic #: 3

[\[All DP-300 Questions\]](#)

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.

```
"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,
            ...
          }
        }
      ]
    }
  }
]
```

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 37

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure SQL managed instance named MI1.

You need to implement automatic tuning for the databases of MI1.

What should you do?

- A. Use The REST API to call the patch operation and modify the AutomaticTuningServerMode property.
- B. From the Azure portal, configure automatic tuning.
- C. Use Transact-SQL to enable the FORCE_LAST_GOOD_PLAN option.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 38

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure subscription that contains an Azure SQL database named db1.

You need to implement SQL insights for db1.

Which two resources should you create first? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. a storage account
- B. a virtual machine
- C. an Azure logic app
- D. an Azure function
- E. a Log Analytics workspace

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 39

Topic #: 3

[\[All DP-300 Questions\]](#)

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

Name	Type
App1	Azure web app
db1	Azure SQL database in the serverless tier

App1 experiences transient connection errors and timeouts when it attempts to access db1 after extended periods of inactivity.

You need to modify db1 to resolve the issues experienced by App1 as soon as possible, without considering immediate costs.

What should you do?

- A. Enable automatic tuning for db1.
- B. Increase the number of vCores allocated to db1.
- C. Decrease the auto-pause delay for db1.
- D. Disable auto-pause delay for db1.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 40

Topic #: 3

[\[All DP-300 Questions\]](#)


HOTSPOT

-

You have an Azure SQL database named DB1 that contains a table named Table1.

You run a query to load data into Table1.

The performance metrics of Table1 during the load operation are 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.

Answer Area

To reduce how long it takes to complete the query you must **[answer choice]**.

- scale the resource
- use an elastic resource
- perform query tuning

To reduce the log IO load of the operation, the query must be updated to use **[answer choice]** table.

- a temporary
- an In-Memory OTLP durable
- an In-Memory OTLP non durable

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 41

Topic #: 3

[\[All DP-300 Questions\]](#)

DRAG DROP

You have a database named db1.

The log for db1 contains the following entry.

```
Date 10/5/2021 10:57:08 AM
```

```
Log SQL Server (Current - 10/5/2021 11:26:00 AM)
```

```
Sources spid1595
```

```
Message
```

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

You need to ensure that db1 can process transactions.

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

Back up the transaction log file.

Remove db1 from the availability group.

Shrink the transaction log file.

Shrink db1.

Add db1 back to the availability group.

Answer Area

1

2

3



Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 42

Topic #: 3

[\[All DP-300 Questions\]](#)

You have the following resources:

- 15 SQL Server on Azure Virtual Machines instances
- 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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 43

Topic #: 3

[\[All DP-300 Questions\]](#)

HOTSPOT

-

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

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

- Create jobs.
- View all jobs.
- Modify, delete, and disable the jobs the user created.

The solution must use the principle of least privilege.

Which built-in database role should you assign to User1, and where is the role defined? To answer, select the appropriate options in the answer area.

Answer Area

Location:

DB1
master
msdb

Built-in role:

SQLAgentOperatorRole
SQLAgentReaderRole
SQLAgentUserRole

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 44

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a Log Analytics workspace named Workspace1.

You need to collect performance metrics for SQLMI1 and stream the metrics to Workspace.

What should you do first?

- A. Create a private endpoint connection on SQLMI1.
- B. Configure Azure SQL Analytics to use Workspace1.
- C. Modify the Compute + storage settings for SQLMI1.
- D. Modify the diagnostic settings for SQLMI1.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 45

Topic #: 3

[\[All DP-300 Questions\]](#)

HOTSPOT

-

You have an Azure SQL database named DB1 in the General Purpose service tier.

You need to monitor DB1 by using SQL Insights.

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

Answer Area

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	

Show Suggested Answer

Actual exam question from Microsoft's DP-300


Question #: 46

Topic #: 3

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 47

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure SQL database named DB1 that contains a nonclustered index named index1.

End users report slow queries when they use index1.

You need to identify the operations that are being performed on the index.

Which dynamic management view should you use?

- A. Sys.dm_exec_query_plan_stats
- B. Sys.dm_db_index_physical_stats
- C. Sys.dm_db_index_operational_stats
- D. Sys.dm_db_index_usage_stats

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 48

Topic #: 3

[\[All DP-300 Questions\]](#)

HOTSPOT

-

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

You need to monitor the performance of SQLMI1 and identify which database uses the most memory and the most disk I/O.

Which objects should you query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

For memory:

 sys.dm_exec_query_stats
 sys.dm_io_virtual_file_stats
 sys.dm_os_buffer_descriptors
 sys.dm_os_performance_counters

For disk I/O:

 sys.dm_exec_query_stats
 sys.dm_io_virtual_file_stats
 sys.dm_os_buffer_descriptors
 sys.dm_os_performance_counters

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 49

Topic #: 3

[\[All DP-300 Questions\]](#)

HOTSPOT

-

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines. The virtual machine hosts a database named DB1.

You need to monitor DB1 by using Extended Events. The solution must meet the following requirements:

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

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

NOTE: Each correct selection is worth one point.

Answer Area

```
CREATE EVENT SESSION session1 ON DATABASE
```

```
ADD EVENT sqlserver.sql_statement_starting
```

```
(
```

```
  ACTION (sqlserver.sql_text)
```

```
  WHERE statement LIKE 'UPDATE gmTabEmployee%'
```

```
)
```

```
ADD TARGET
```

```
package0.
```

event_file
event_stream
ring_buffer

```
(
```

```
  SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'
```

```
)
```

```
WITH
```

```
(MAX_MEMORY = 10 MB,
```

```
  EVENT_RETENTION_MODE=
```

ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_SINGLE_EVENT_LOSS
NO_EVENT_LOSS

```
  MAX_DISPATCH_LATENCY = 3 SECONDS)
```

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 50

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure SQL database named DB1.

You need to query the fragmentation information of data and indexes for the tables in DB1.

Which command should you run?

- A. `sys.dm_db_index_usage_stats`
- B. `DBCC CHECKALLOC`
- C. `DBCC SHOWCONTIG`
- D. `sts.dm_db_index_physical_stats`

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 51

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure SQL managed instance named SQLMI1 that has the following settings:

- vCores: 4
- Service tier: General Purpose
- Hardware generation: Standard-series (Gen5)

You discover that memory pressure on SQLMI1 is high.

You need to reduce the memory pressure on SQLMI1. The solution must minimize costs.

What should to do?

- A. Enable the Query Store.
- B. Change vCores to 8.
- C. Change Hardware generation to Premium-series.
- D. Change Service tier to Business Critical.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 52

Topic #: 3

[\[All DP-300 Questions\]](#)

You have five instances of SQL Server on Azure Virtual Machines.

You need to monitor Microsoft SQL Server performance for all the instances by consolidating metrics into a single graphic display. The solution must minimize administrative effort.

What should you use?

- A. Azure Monitor
- B. Log Analytics
- C. SQL Insights
- D. Azure SQL Analytics

[Show Suggested Answer](#)





Actual exam question from Microsoft's DP-300

Question #: 53

Topic #: 3

[\[All DP-300 Questions\]](#)

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

SQL1 contains an Extended Events session named session1 that captures Microsoft SQL Server events.

You need to correlate the session events with events captured by Event Tracing for Windows (ETW).

What should you do for session1?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 54

Topic #: 3

[\[All DP-300 Questions\]](#)

You have an Azure subscription that contains the following resources:

- 10 Azure SQL databases
- Five Azure SQL managed instances
- Five instances of SQL Server on Azure Virtual Machines

You need to implement a centralized monitoring solution for all the Azure SQL resources. The solution must minimize administrative effort.

What should you include in the solution?

- A. Log Analytics
- B. Azure SQL Analytics
- C. Query Performance Insight
- D. SQL Insights

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 59

Topic #: 3

[\[All DP-300 Questions\]](#)

Case study -

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview -

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment. Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Existing Environment. Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Existing Environment. Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Existing Environment. Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements. Planned Changes -

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Requirements. Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Requirements. Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Requirements. Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Requirements. Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to recommend a solution to meet the security requirements and the business requirements for DB3.

What should you recommend as the first step of the solution?

- Run the sp_addarticle stored procedure.
- Run the ALTER TABLE statement and specify the ENABLE CHANGE_TRACKING Clause.
- Run the ALTER DATABASE statement and specify the SET CHANGE_TRACKING = ON Clause.
- Run the sys.sp_cdc_enable_db stored procedure.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 60

Topic #: 3

[\[All DP-300 Questions\]](#)

Case study -

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview -

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment. Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Existing Environment. Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Existing Environment. Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Existing Environment. Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements. Planned Changes -

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Requirements. Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Requirements. Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Requirements. Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Requirements. Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to recommend a solution to ensure that the performance of DB3 is optimized after the migration to Azure SQL Database. The solution must meet availability requirements.

What should you include in the recommendation?

- A. vertical scaling
- B. a custom resource pool
- C. Resource Governor
- D. horizontal scaling

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 61

Topic #: 3

[\[All DP-300 Questions\]](#)

Case study -

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview -

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment. Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Existing Environment. Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Existing Environment. Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Existing Environment. Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements. Planned Changes -

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Requirements. Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Requirements. Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Requirements. Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Requirements. Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to identify the event_file target for monitoring DB3 after the migration to Azure SQL Database. The solution must meet the management requirements.

What should you use as the event_file target?

- A. a SQL Server filegroup
- B. an Azure SQL database
- C. an Azure Files share
- D. an Azure Blob Storage container

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 1

Topic #: 4

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 2

Topic #: 4

[\[All DP-300 Questions\]](#)

HOTSPOT -

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.

Hot Area:

Answer Area

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 3

Topic #: 4

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 4

Topic #: 4

[\[All DP-300 Questions\]](#)

You have an Azure SQL database.

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

You run the `select * from sys.database_scoped_configurations` Transact-SQL command and receive the results shown in the following table.

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

You need relieve the memory pressure.

What should you configure?

- A. LEGACY_CARDINALITY_ESTIMATION
- B. QUERY_OPTIMIZER_HOTFIXES
- C. OPTIMIZE_FOR_AD_HOC_WORKLOADS
- D. ACCELERATED_PLAN_FORCING

Show Suggested Answer

Actual exam question from Microsoft's DP-300

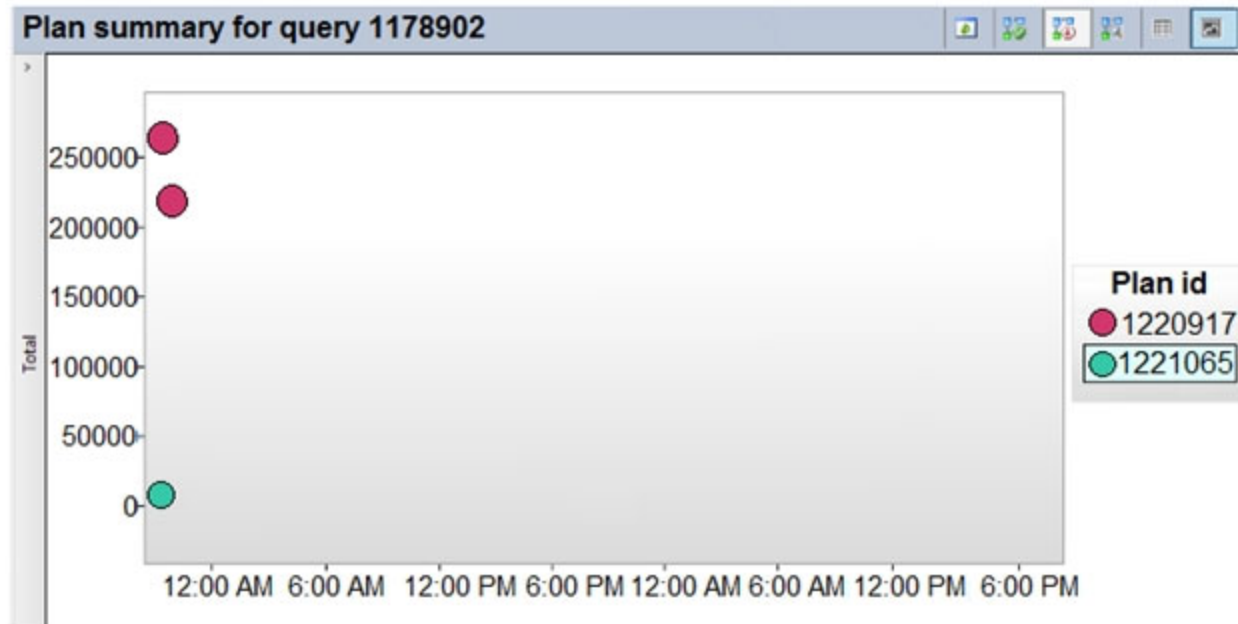
Question #: 5

Topic #: 4

[\[All DP-300 Questions\]](#)

You have SQL Server on an Azure virtual machine that contains a database named DB1.

You view a plan summary that shows the duration in milliseconds of each execution of query 1178902 as shown in the following exhibit:



What should you do to ensure that the query uses the execution plan which executes in the least amount of time?

- A. Force the query execution plan for plan 1221065.
- B. Run the DBCC FREEPROCCACHE command.
- C. Force the query execution plan for plan 1220917.
- D. Disable parameter sniffing.

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 6

Topic #: 4

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.


 Azure SQL Database built-in intelligence automatically tunes your databases to optimize performance. Click here to learn more about automatic tuning. 

Inherit from: 




Server

Azure defaults

Don't inherit

 The database is inheriting automatic tuning configuration from Azure defaults.

Configure the automatic tuning options 

OPTION	DESIRED STATE	CURRENT STATE
 FORCE PLAN	<input type="radio"/> ON <input type="radio"/> OFF <input checked="" type="radio"/> INHERIT	ON Auto-configured by Azure
 CREATE INDEX	<input type="radio"/> ON <input type="radio"/> OFF <input checked="" type="radio"/> INHERIT	ON Auto-configured by Azure
 DROP INDEX	<input checked="" type="radio"/> ON <input type="radio"/> OFF <input type="radio"/> INHERIT	ON Forced by user

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
Nonclustered indexes will be added to tables to improve performance.	<input type="radio"/>	<input type="radio"/>
Columns will be added to existing indexes automatically.	<input type="radio"/>	<input type="radio"/>
The query execution plan will revert to a previous plan if query performance degrades.	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 7

Topic #: 4

[\[All DP-300 Questions\]](#)

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

You review the actual execution plan for the query, and you add an index to a table referenced by the query.

You need to compare the previous actual execution plan for the query to the Live Query Statistics.

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

- A. For DB1, set QUERY_CAPTURE_MODE of Query Store to All.
- B. Run the SET SHOWPLAN_ALL Transact-SQL statement.
- C. Save the actual execution plan.
- D. Enable Query Store for DB1.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 8

Topic #: 4

[\[All DP-300 Questions\]](#)

You have an Azure SQL database.

Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue.

You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort.

What should you use?

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 9

Topic #: 4

[\[All DP-300 Questions\]](#)

You have an Azure SQL database. The database contains a table that uses a columnstore index and is accessed infrequently.

You enable columnstore archival compression.

What are two possible results of the configuration? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Queries that use the index will consume more disk I/O.
- B. Queries that use the index will retrieve fewer data pages.
- C. The index will consume more disk space.
- D. The index will consume more memory.
- E. Queries that use the index will consume more CPU resources.

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 10

Topic #: 4

[\[All DP-300 Questions\]](#)

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

[Show Suggested Answer](#)



Actual exam question from Microsoft's DP-300

Question #: 11

Topic #: 4

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 12

Topic #: 4

[\[All DP-300 Questions\]](#)

HOTSPOT -

You are designing an enterprise data warehouse in Azure Synapse Analytics that will store website traffic analytics in a star schema.

You plan to have a fact table for website visits. The table will be approximately 5 GB.

You need to recommend which distribution type and index type to use for the table. The solution must provide the fastest query performance.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Distribution:

	▼
Hash	
Round robin	
Replicated	

Index:

	▼
Clustered columnstore	
Clustered	
Nonclustered	

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 13

Topic #: 4

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 14

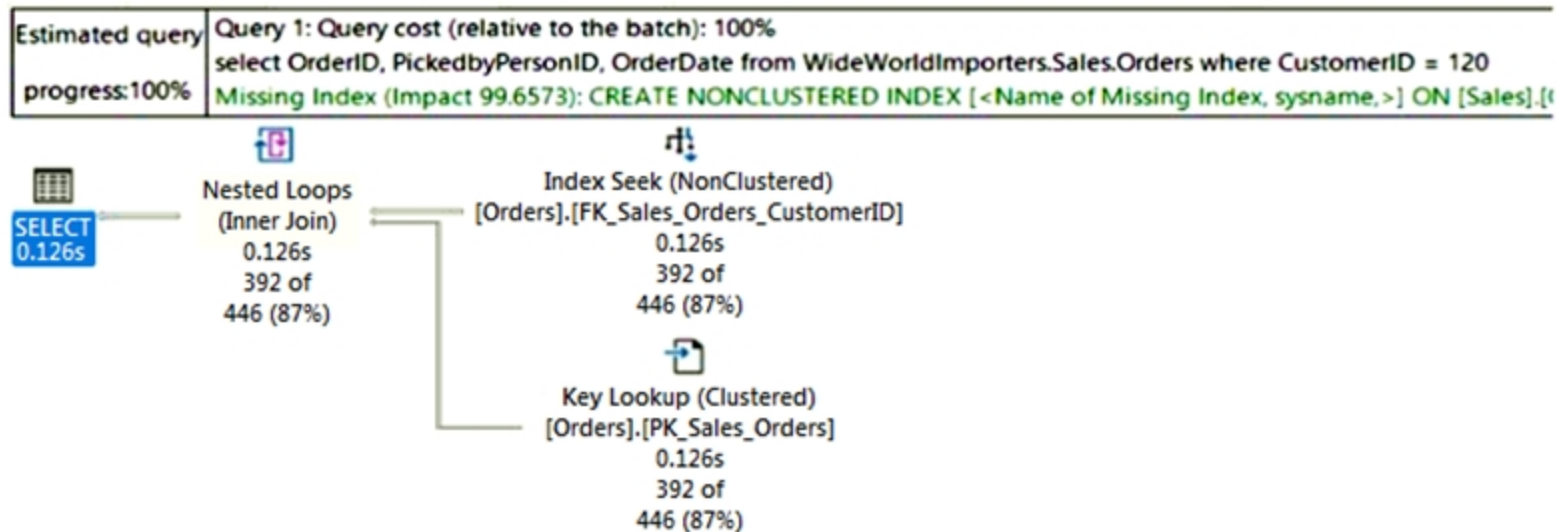
Topic #: 4

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have an Azure SQL database.

You are reviewing a slow performing query 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.

Hot Area:

Answer Area

The exhibit shows **[answer choice]**.

- an actual execution plan
- an estimated execution plan
- Live Query Statistics

The **[answer choice]** operator in the execution plan indicates that the query would benefit from performance tuning.

- Index Seek
- Key Lookup
- Nested Loops

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 15

Topic #: 4

[\[All DP-300 Questions\]](#)

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_stats`.
- D. Generate Live Query Statistics.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 16

Topic #: 4

[\[All DP-300 Questions\]](#)

HOTSPOT -

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

You use Query Performance Insight to monitor db1.

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Configuration setting:

- DATA_FLUSH_INTERVAL_SECONDS
- INTERVAL_LENGTH_MINUTES
- MAX_PLANS_PER_QUERY
- QUERY_CAPTURE_MODE

Value:

- 1
- 60
- CUSTOM
- ON

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 17

Topic #: 4

[\[All DP-300 Questions\]](#)

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. SQL Profiler
- B. Extended Events
- C. Query Store
- D. dynamic management views

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 18

Topic #: 4

[\[All DP-300 Questions\]](#)

You have an Azure SQL database named DB1.

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

What should you do first?

- A. Run the SET SHOWPLAN_ALL Transact-SQL statement.
- B. For DB1, set QUERY_CAPTURE_MODE of Query Store to All.
- C. Run the SET FORCEPLAN Transact-SQL statement.
- D. Enable Query Store for DB1.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 19

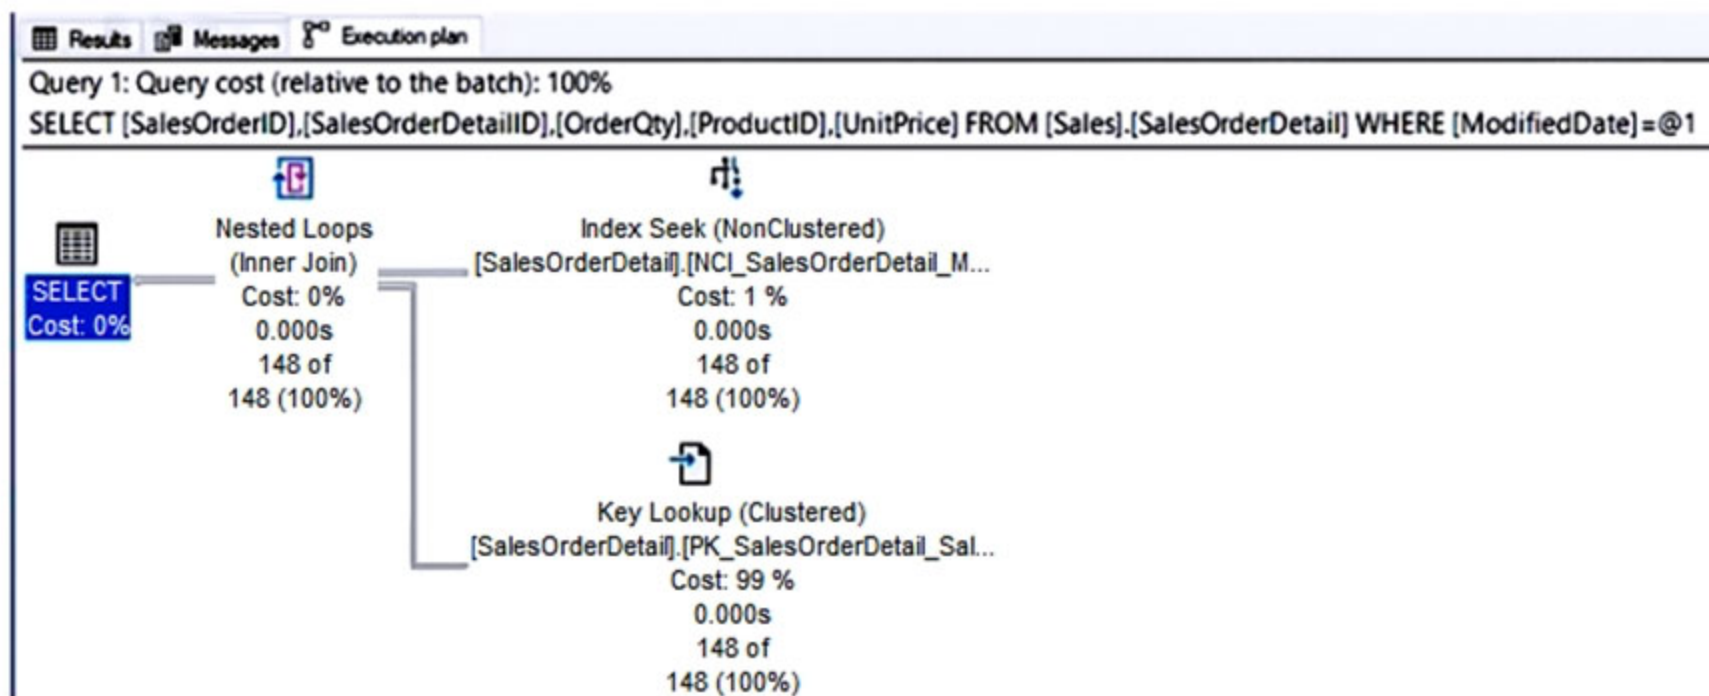
Topic #: 4

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have an Azure SQL database.

You have a query and the associated execution plan 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.

Hot Area:

Answer Area

The performance issue stems from the **[answer choice]** operator.

- Select
- Index Seek
- Key Lookup
- Nested Loops

The performance issue can be resolved by adding include columns to the **[answer choice]**.

- heap
- clustered index
- nonclustered index

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 20

Topic #: 4

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 21

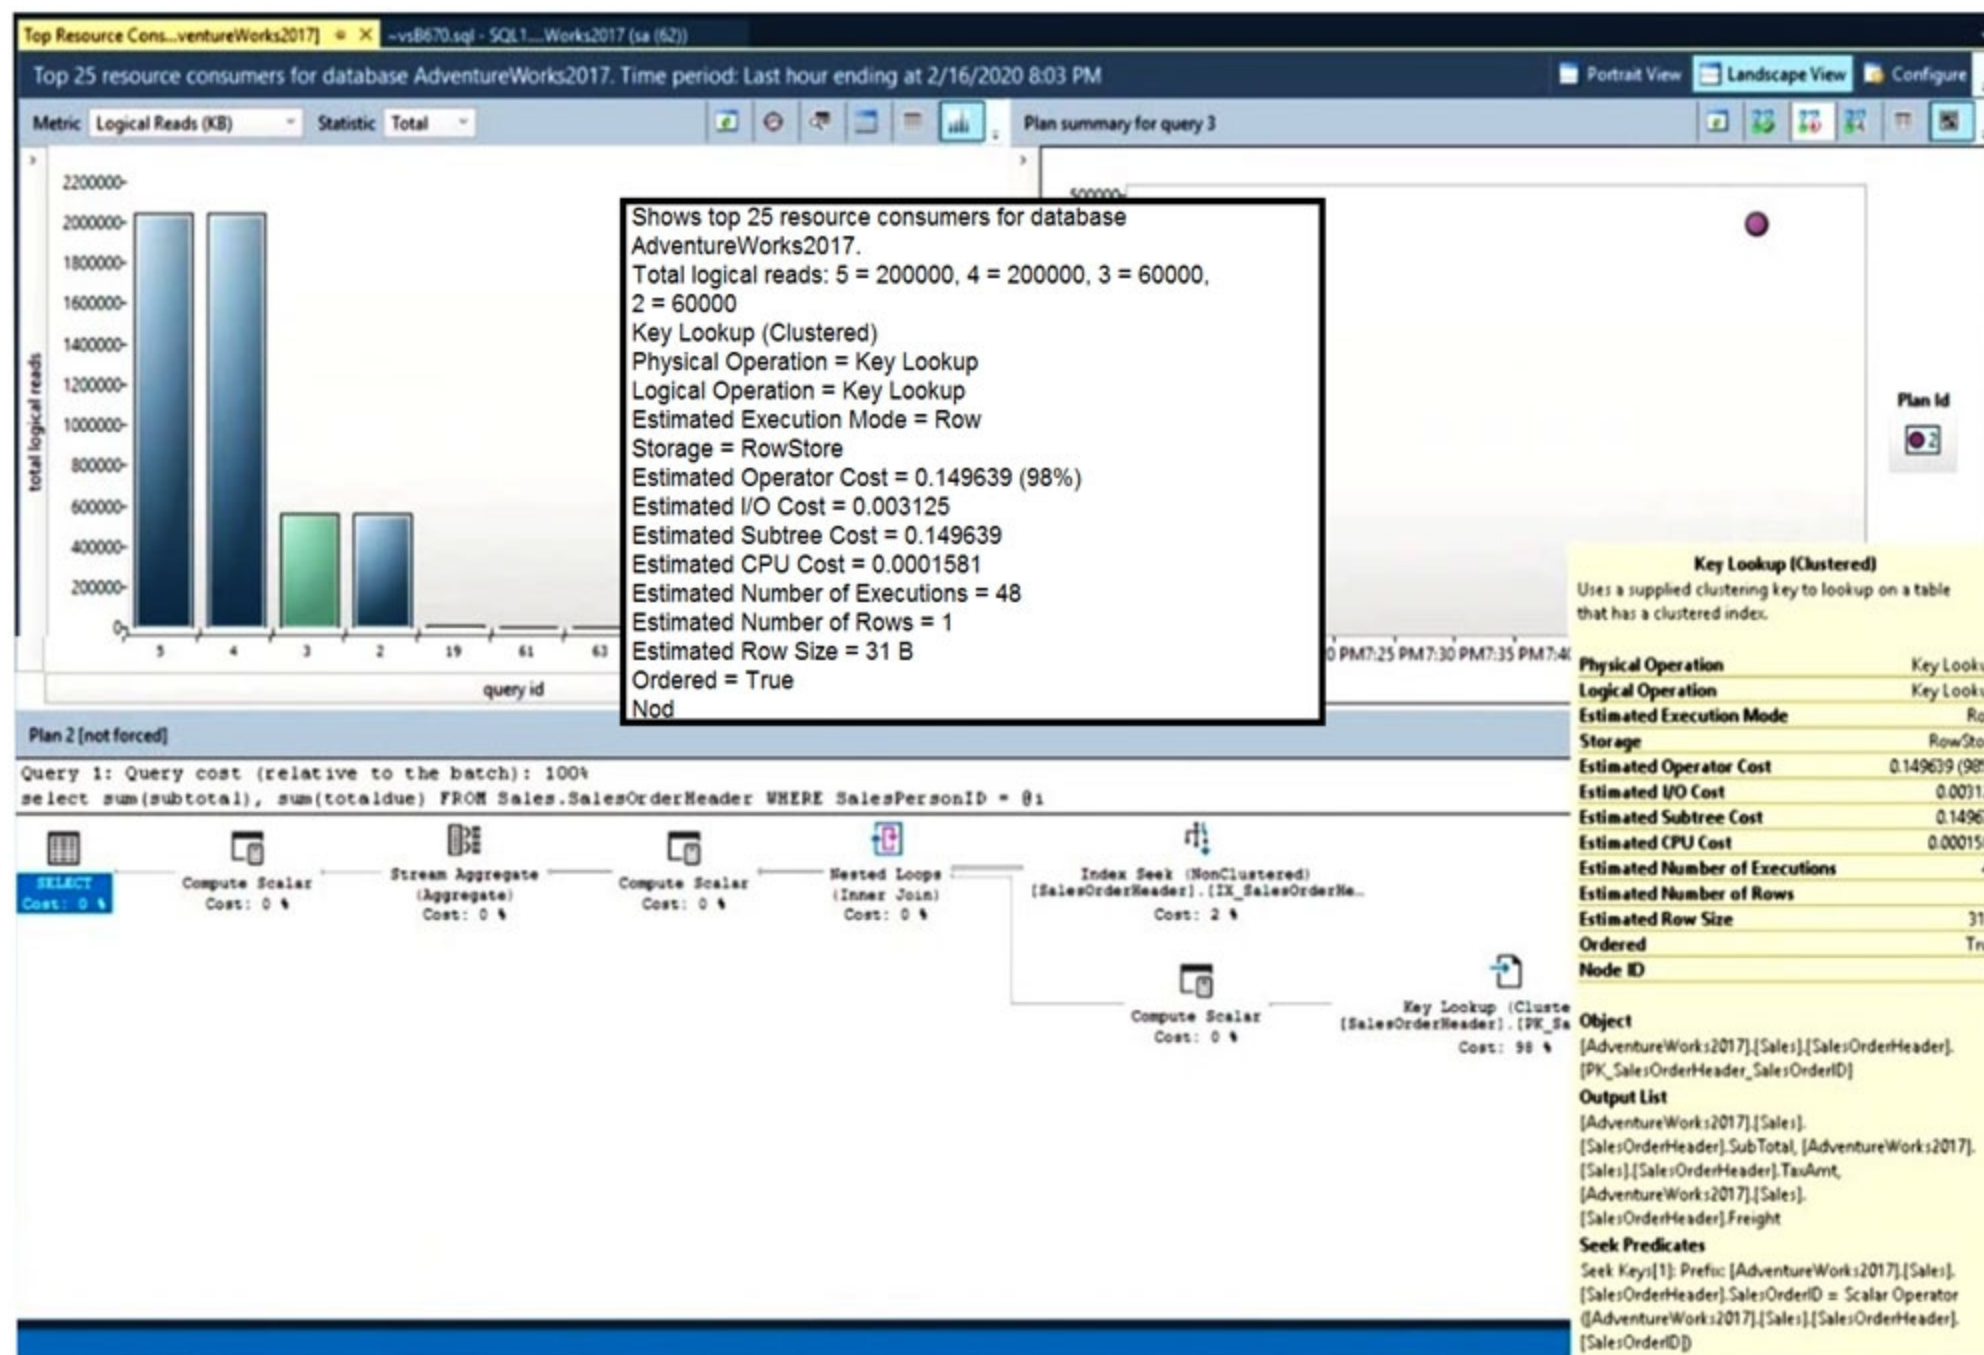
Topic #: 4

[\[All DP-300 Questions\]](#)

HOTSPOT -

You have SQL Server on an Azure virtual machine.

You review the query plan shown in the following exhibit.



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

You will reduce the I/O usage and the query execution time if you force the query plan.

Yes

No

You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.

You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 22

Topic #: 4

[\[All DP-300 Questions\]](#)

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.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 23

Topic #: 4

[\[All DP-300 Questions\]](#)

You have an Azure SQL database.

You identify a long running query.

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

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

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 24

Topic #: 4

[\[All DP-300 Questions\]](#)

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

You need to identify which database queries consume the most resources.

Which tool should you use?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 25

Topic #: 4

[\[All DP-300 Questions\]](#)

You have an Azure subscription that contains an Azure SQL database. The database contains a table named table1.

You execute the following Transact-SQL statements.

```
CREATE CLUSTERED INDEX PK_index1 ON [dbo].[table1] ([Column1])  
CREATE NONCLUSTERED INDEX [NCI_index2] ON [dbo].[table1] ([Column1]) ASC
```

You need to reduce the time it takes to perform analytic queries on the database.

Which configuration should you enable?

- A. ROW_MODE_MEMORY_GRANT_FEEDBACK
- B. BATCH_MODE_MEMORY_GRANT_FEEDBACK
- C. BATCH_MODE_ADAPTIVE_JOINS
- D. BATCH_MODE_ON_ROWSTORE

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 26

Topic #: 4

[\[All DP-300 Questions\]](#)

DRAG DROP

You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored procedures.

You need to ensure that SQL Server Agent jobs running on SQL1 can notify administrators when a failure occurs.

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

Create a Database Mail account.

Create a profile named
AzureManagedInstance_dbmail_profile.

Create a profile named
application_dbmail_profile.

Enable email notifications upon failure.

Enable pager notifications upon failure.

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 27

Topic #: 4

[\[All DP-300 Questions\]](#)

You have a Microsoft SQL Server 2019 database named DB1 and an Azure SQL managed instance named SQLMI1.

You need to move a SQL Server Agent job from DB1 to SQLMI1.

Which job attribute is unsupported in SQLMI1?

- A. log to table
- B. email notifications
- C. schedules
- D. output files

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 28

Topic #: 4

[\[All DP-300 Questions\]](#)

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.

Actions Commands Cmdlets Statements

Configure Database Mail.

Configure a job notification.

Configure an alert.

Configure SendGrid.

Create an operator.



Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 29

Topic #: 4

[\[All DP-300 Questions\]](#)

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';

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 1

Topic #: 5

[\[All DP-300 Questions\]](#)

HOTSPOT -

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.

Hot Area:

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.

Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 2

Topic #: 5

[\[All DP-300 Questions\]](#)

You have the following Azure Data Factory pipelines:

⇒ Ingest Data from System1

Ingest Data from System2 -

-
- ⇒ Populate Dimensions
- ⇒ Populate Facts

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

What should you do to schedule the pipelines for execution?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 3

Topic #: 5

[\[All DP-300 Questions\]](#)

You have an Azure Data Factory pipeline that performs an incremental load of source data to an Azure Data Lake Storage Gen2 account.

Data to be loaded is identified by a column named LastUpdatedDate in the source table.

You plan to execute the pipeline every four hours.

You need to ensure that the pipeline execution meets the following requirements:

- ⇒ Automatically retries the execution when the pipeline run fails due to concurrency or throttling limits.
- ⇒ Supports backfilling existing data in the table.

Which type of trigger should you use?

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 4

Topic #: 5

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 5


Topic #: 5

[\[All DP-300 Questions\]](#)


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:

- ⇒ 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.

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"/>

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 6

Topic #: 5

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 7

Topic #: 5

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 8

Topic #: 5

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 9

Topic #: 5

[\[All DP-300 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Data Lake Storage account that contains a staging zone.

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

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

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 10

Topic #: 5

[\[All DP-300 Questions\]](#)

DRAG DROP -

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.

Select and Place:

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



Show Suggested Answer

Actual exam question from Microsoft's DP-300

Question #: 11

Topic #: 5

[\[All DP-300 Questions\]](#)

DRAG DROP -

You have SQL Server on an Azure virtual machine.

You need to use Policy-Based Management in Microsoft SQL Server to identify stored procedures that do not comply with your naming conventions.

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

Export a built-in policy.

Create a custom policy based on a condition.

Create a custom condition based on a built-in facet.

View the policy history.

Import a policy file.

Run a policy evaluation.

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 12

Topic #: 5

[\[All DP-300 Questions\]](#)

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

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

- ⇒ Minimize costs.
- ⇒ Aggregate Intelligent Insights telemetry from each database.

What should you do?

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

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 13

Topic #: 5

[\[All DP-300 Questions\]](#)

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

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 14

Topic #: 5

[\[All DP-300 Questions\]](#)

You have an Azure SQL managed instance.

You need to enable SQL Agent Job email notifications.

What should you do?

- A. Use the Agent XPs option.
- B. Enable the SQL Server Agent.
- C. Run the `sp_configure` command.
- D. Run the `sp_set_agent_properties` command.

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 15

Topic #: 5

[\[All DP-300 Questions\]](#)

You have four Azure subscriptions. Each subscription contains multiple Azure SQL databases.

You need to update the column and index statistics for the databases.

What should you use?

- A. an Azure Automation runbook
- B. a SQL Agent job
- C. Azure SQL Analytics
- D. automatic tuning in Azure SQL Database

Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 16

Topic #: 5

[\[All DP-300 Questions\]](#)

DRAG DROP -

You have SQL Server on an Azure virtual machine named SQL1.

SQL1 has an agent job to back up all databases.

You add a user named dbadmin1 as a SQL Server Agent operator.

You need to ensure that dbadmin1 receives an email alert if a 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.

Select and Place:

Actions

Answer Area

Create a job alert

Create a job notification

Enable Database Mail

Enable the email settings for the SQL Server Agent

Create a job target



Show Suggested Answer



Actual exam question from Microsoft's DP-300

Question #: 17

Topic #: 5

[\[All DP-300 Questions\]](#)

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



Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 18

Topic #: 5

[\[All DP-300 Questions\]](#)

You have an Azure SQL Database managed instance named SQLMI1. A Microsoft SQL Server Agent job runs on SQLMI1.

You need to ensure that an automatic email notification is sent once the job completes.

What should you include in the solution?

- A. From SQL Server Configuration Manager (SSCM), enable SQL Server Agent
- B. From SQL Server Management Studio (SSMS), run `sp_set_sqlagent_properties`
- C. From SQL Server Management Studio (SSMS), create a Database Mail profile
- D. From the Azure portal, create an Azure Monitor action group that has an Email/SMS/Push/Voice action

Show Suggested Answer





Actual exam question from Microsoft's DP-300

Question #: 19

Topic #: 5

[\[All DP-300 Questions\]](#)

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

[Show Suggested Answer](#)





Actual exam question from Microsoft's DP-300

Question #: 20

Topic #: 5

[\[All DP-300 Questions\]](#)

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

You need to create a SQL Server Agent job that will rebuild indexes of the databases hosted on VM1. The solution must use the principle of least privilege.

What should you create first?

- A. a local Windows account
- B. a user-assigned managed identity in Azure AD
- C. a system-assigned managed identity in Azure AD
- D. an Elastic Job agent

Show Suggested Answer

