



DP-300^{Q&As}

Administering Relational Databases on Microsoft Azure

Pass Microsoft DP-300 Exam with 100% Guarantee

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

<https://www.lead4pass.com/dp-300.html>

100% Passing Guarantee
100% Money Back Assurance

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

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





QUESTION 1

You have the following Azure Data Factory pipelines:

1.
Ingest Data from System1
2.
Ingest Data from System2
3.
Populate Dimensions
4.
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.

Correct Answer: D

Reference: <https://www.mssqltips.com/sqlservertip/6137/azure-data-factory-control-flow-activities-overview/>

QUESTION 2

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.

Correct Answer: B

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools. Serverless SQL pool enables you to query data

in your data lake. It offers a T-SQL query surface area that accommodates semi-structured and unstructured data queries.

To support a smooth experience for in place querying of data that's located in Azure Storage files, serverless SQL pool uses the OPENROWSET function with additional capabilities.

The easiest way to see to the content of your JSON file is to provide the file URL to the OPENROWSET function, specify csv FORMAT.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-json-files>

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-data-storage>

QUESTION 3

You plan to move two 100-GB databases to Azure.

You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.

What should you use?

A. two Azure SQL Databases in an elastic pool

B. two databases hosted in SQL Server on an Azure virtual machine

C. two databases in an Azure SQL Managed instance

D. two single Azure SQL databases

Correct Answer: A

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

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

QUESTION 4



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

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

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

1.

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

2.

Minimize the possibility of VM1 exfiltrating data stored in SqlDb1. What should you create on VNet1?

A. a VPN gateway

B. a service endpoint

C. a private link

D. an ExpressRoute gateway

Correct Answer: C

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

Reference:

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

QUESTION 5

HOTSPOT

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

Type	Designated retention period
Application	360 days
Infrastructure	60 days

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

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

1.

Automatically deletes the logs at the end of each retention period



2.
Minimizes storage costs
What should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

To minimize storage costs:

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

To delete the logs automatically:

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

Correct Answer:

Answer Area

To minimize storage costs:

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

To delete the logs automatically:

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

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

Hot - Optimized for storing data that is accessed frequently.

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

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

Box 2: Azure Blob storage lifecycle management rules

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

Reference:

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



QUESTION 6

You receive numerous alerts from Azure Monitor for an Azure SQL database.

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

Correct Answer: BD

B: Dynamic Thresholds continuously learns the data of the metric series and tries to model it using a set of algorithms and methods. It detects patterns in the data such as seasonality (Hourly / Daily / Weekly), and is able to handle noisy metrics (such as machine CPU or memory) as well as metrics with low dispersion (such as availability and error rate).

D: Alert threshold sensitivity is a high-level concept that controls the amount of deviation from metric behavior required to trigger an alert.

Low - The thresholds will be loose with more distance from metric series pattern. An alert rule will only trigger on large deviations, resulting in fewer alerts.

Incorrect Answers:

A: High - The thresholds will be tight and close to the metric series pattern. An alert rule will be triggered on the smallest deviation, resulting in more alerts.

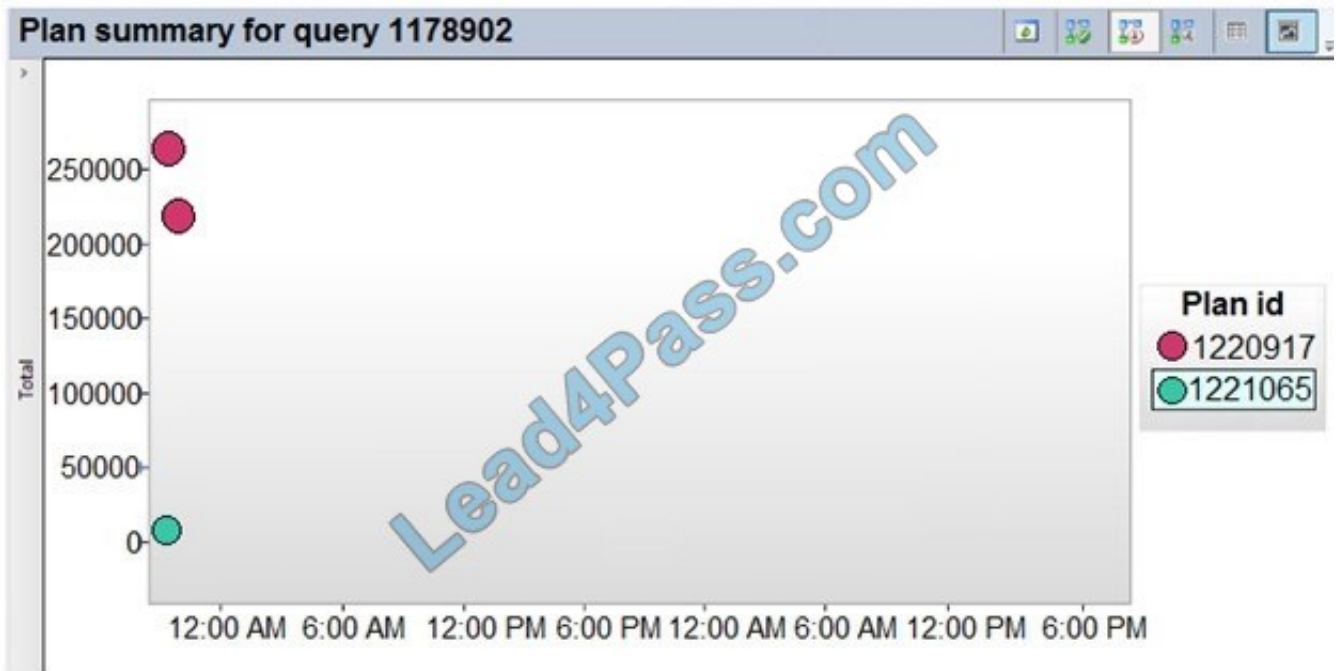
Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-dynamic-thresholds>

QUESTION 7

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.

Correct Answer: C

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-store-usage-scenarios>

QUESTION 8

Based on the PaaS prototype, which Azure SQL Database compute tier should you use?

- A. Business Critical 4-vCore
- B. Hyperscale
- C. General Purpose v-vCore
- D. Serverless

Correct Answer: A

There are CPU and Data I/O spikes for the PaaS prototype. Business Critical 4-vCore is needed.

Incorrect Answers:



B: Hyperscale is for large databases

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

QUESTION 9

HOTSPOT

You have SQL Server on an Azure virtual machine that contains a database named DB1.

The database reports a CHECKSUM error.

You need to recover the database.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

USE master;

ALTER DATABASE [DB1] SET

	▼
OFFLINE	
ONLINE	
SINGLE_USER	
TRUSTWORTHY	

 WITH ROLLBACK IMMEDIATE;

GO

DBCC CHECKDB ('DB1',

	▼
MOINDEX	
PHYSICAL_ONLY	
REPAIR_ALLOW_DATA_LOSS	
REPAIR_FAST	

 WITH NO_INFOMSGS;

GO

ALTER DATABASE [DB1] SET

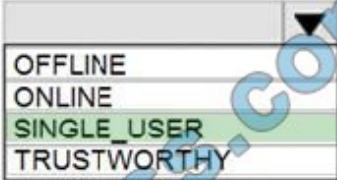
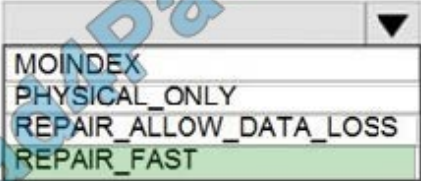
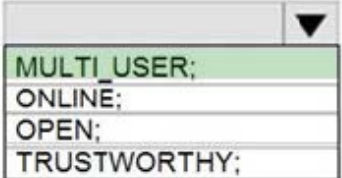
	▼
MULTI_USER;	
ONLINE;	
OPEN;	
TRUSTWORTHY;	

GO

Correct Answer:



Answer Area

```
USE master;  
ALTER DATABASE [DB1] SET  WITH ROLLBACK IMMEDIATE;  
GO  
  
DBCC CHECKDB ('DB1',  WITH NO_INFOMSGS;  
GO  
  
ALTER DATABASE [DB1] SET   
GO
```

Box 1: SINGLE_USER

The specified database must be in single-user mode to use one of the following repair options.

Box 2: REPAIR_ALLOW_DATA_LOSS

REPAIR_ALLOW_DATA_LOSS tries to repair all reported errors. These repairs can cause some data loss.

Note: The REPAIR_ALLOW_DATA_LOSS option is a supported feature but it may not always be the best option for bringing a database to a physically consistent state. If successful, the REPAIR_ALLOW_DATA_LOSS option may result in

some data loss. In fact, it may result in more data lost than if a user were to restore the database from the last known good backup.

Incorrect Answers:

REPAIR_FAST

Maintains syntax for backward compatibility only. No repair actions are performed.

Box 3: MULTI_USER

MULTI_USER

All users that have the appropriate permissions to connect to the database are allowed.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql>



QUESTION 10

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="checkbox"/> ON <input type="checkbox"/> OFF <input checked="" type="checkbox"/> INHERIT	ON Auto-configured by Azure
CREATE INDEX	<input type="checkbox"/> ON <input type="checkbox"/> OFF <input checked="" type="checkbox"/> INHERIT	ON Auto-configured by Azure
DROP INDEX	<input checked="" type="checkbox"/> ON <input type="checkbox"/> OFF <input type="checkbox"/> 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"/>

Correct Answer:

Answer Area

Statements	Yes	No
Nonclustered indexes will be added to tables to improve performance.	<input checked="" type="radio"/>	<input type="radio"/>
Columns will be added to existing indexes automatically.	<input type="radio"/>	<input checked="" type="radio"/>
The query execution plan will revert to a previous plan if query performance degrades.	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: Yes

We see: Tuning option: Create index ON

CREATE INDEX - Identifies indexes that may improve performance of your workload, creates indexes, and automatically verifies that performance of queries has improved.

Box 2: No

Box 3: Yes

FORCE LAST GOOD PLAN (automatic plan correction) - Identifies Azure SQL queries using an execution plan that is slower than the previous good plan, and queries using the last known good plan instead of the regressed plan.



QUESTION 11

You need to implement the surrogate key for the retail store table. The solution must meet the sales transaction dataset requirements.

What should you create?

- A. a table that has a FOREIGN KEY constraint
- B. a table that has an IDENTITY property
- C. a user-defined SEQUENCE object
- D. a system-versioned temporal table

Correct Answer: B

Scenario: Contoso requirements for the sales transaction dataset include: Implement a surrogate key to account for changes to the retail store addresses.

A surrogate key on a table is a column with a unique identifier for each row. The key is not generated from the table data. Data modelers like to create surrogate keys on their tables when they design data warehouse models. You can use the IDENTITY property to achieve this goal simply and effectively without affecting load performance.

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

QUESTION 12

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;

Correct Answer:

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;



Box 1: sys.resource_stats

sys.resource_stats returns CPU usage and storage data for an Azure SQL Database. It has database_name and start_time columns.

Box 2: DateAdd

The following example returns all databases that are averaging at least 80% of compute utilization over the last one week.

```
DECLARE @s datetime;
```

```
DECLARE @e datetime;
```

```
SET @s= DateAdd(d,-7,GetUTCDate());
```

```
SET @e= GETUTCDATE();
```

```
SELECT database_name, AVG(avg_cpu_percent) AS Average_Compute_Utilization FROM sys.resource_stats  
WHERE start_time BETWEEN @s AND @e GROUP BY database_name HAVING AVG(avg_cpu_percent) >= 80
```

Incorrect Answers: sys.dm_exec_requests: sys.dm_exec_requests returns information about each request that is executing in SQL Server. It does not have a column named database_name.

sys.dm_db_resource_stats:

sys.dm_db_resource_stats does not have any start_time column.

Note: sys.dm_db_resource_stats returns CPU, I/O, and memory consumption for an Azure SQL Database database. One row exists for every 15 seconds, even if there is no activity in the database. Historical data is maintained for approximately one hour.

sys.dm_user_db_resource_governance returns actual configuration and capacity settings used by resource governance mechanisms in the current database or elastic pool. It does not have any start_time column.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-resource-stats-azure-sql-database>

QUESTION 13

You plan to perform batch processing in Azure Databricks once daily. Which type of Databricks cluster should you use?

- A. automated
- B. interactive
- C. High Concurrency

Correct Answer: A

Azure Databricks makes a distinction between all-purpose clusters and job clusters. You use all-purpose clusters to analyze data collaboratively using interactive notebooks. You use job clusters to run fast and robust automated jobs.



The Azure Databricks job scheduler creates a job cluster when you run a job on a new job cluster and terminates the cluster when the job is complete.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/clusters>

[DP-300 PDF Dumps](#)

[DP-300 Study Guide](#)

[DP-300 Braindumps](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

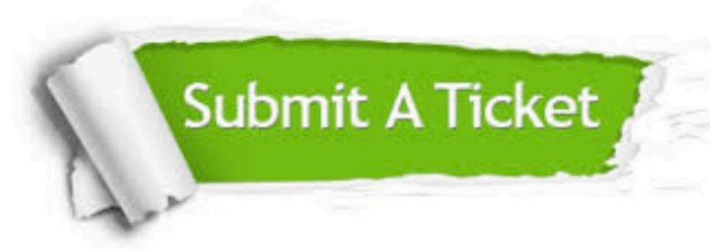
100% Guaranteed Success
100% Money Back Guarantee
365 Days Free Update
Instant Download After Purchase
24x7 Customer Support
Average 99.9% Success Rate
More than 800,000 Satisfied Customers Worldwide
Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.lead4pass.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket:



 <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.
All trademarks are the property of their respective owners.
Copyright © lead4pass, All Rights Reserved.