

## DP-201 Dumps

### Designing an Azure Data Solution

<https://www.certleader.com/DP-201-dumps.html>



**NEW QUESTION 1**

- (Exam Topic 1)

You need to design the Planning Assistance database.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statement	Yes	No
Including a clustered columnstore index in the design will benefit performance.	<input type="radio"/>	<input type="radio"/>
Including a nonclustered columnstore index in the design will benefit performance.	<input type="radio"/>	<input type="radio"/>
Including an index on the License Plate column will benefit performance.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: No

Data used for Planning Assistance must be stored in a sharded Azure SQL Database. Box 2: Yes

Box 3: Yes

Planning Assistance database will include reports tracking the travel of a single vehicle

**NEW QUESTION 2**

- (Exam Topic 1)

You need to ensure that performance requirements for Backtrack reports are met.

What should you recommend? To answer, drag the appropriate technologies to the correct locations. Each technology 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.

Technologies	Answer Area	
	Requirement	Technology
Cosmos DB TTL		
Cosmos DB indexes		
Cosmos DB transactions	Backtrack reporting	<input type="text"/>
Cosmos DB change feed	Privacy and security policy	<input type="text"/>
Cosmos DB stored procedures		

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Cosmos DB indexes

The report for Backtrack must execute as quickly as possible.

You can override the default indexing policy on an Azure Cosmos container, this could be useful if you want to tune the indexing precision to improve the query performance or to reduce the consumed storage.

Box 2: Cosmos DB TTL

This solution reports on all data related to a specific vehicle license plate. The report must use data from the SensorData collection. Users must be able to filter vehicle data in the following ways:

- vehicles on a specific road
- vehicles driving above the speed limit

Note: With Time to Live or TTL, Azure Cosmos DB provides the ability to delete items automatically from a container after a certain time period. By default, you can set time to live at the container level and override the value on a per-item basis. After you set the TTL at a container or at an item level, Azure Cosmos DB will automatically remove these items after the time period, since the time they were last modified.

**NEW QUESTION 3**

- (Exam Topic 1)

You need to design the runtime environment for the Real Time Response system. What should you recommend?

- A. General Purpose nodes without the Enterprise Security package
- B. Memory Optimized Nodes without the Enterprise Security package
- C. Memory Optimized nodes with the Enterprise Security package

D. General Purpose nodes with the Enterprise Security package

**Answer:** B

**NEW QUESTION 4**

- (Exam Topic 2)

You plan to use an Azure SQL data warehouse to store the customer data. You need to recommend a disaster recovery solution for the data warehouse. What should you include in the recommendation?

- A. AzCopy
- B. Read-only replicas
- C. AdICopy
- D. Geo-Redundant backups

**Answer:** D

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore>

**NEW QUESTION 5**

- (Exam Topic 2)

You need to design the encryption strategy for the tagging data and customer data.

What should you recommend? To answer, drag the appropriate setting to the correct drop targets. Each source 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.

Encryption methods	Solution component	Encryption method
Encryption at rest	Tagging data	
Transparent data encryption	Processed customer data	
Azure Key Vault		

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

All cloud data must be encrypted at rest and in transit. Box 1: Transparent data encryption

Encryption of the database file is performed at the page level. The pages in an encrypted database are encrypted before they are written to disk and decrypted when read into memory.

Box 2: Encryption at rest

Encryption at Rest is the encoding (encryption) of data when it is persisted. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption?view= https://docs.microsoft.com/en-us/azure/security/azure-security-encryption-atrest>

**NEW QUESTION 6**

- (Exam Topic 3)

A company stores sensitive information about customers and employees in Azure SQL Database. You need to ensure that the sensitive data remains encrypted in transit and at rest.

What should you recommend?

- A. Transparent Data Encryption
- B. Always Encrypted with secure enclaves
- C. Azure Disk Encryption
- D. SQL Server AlwaysOn

**Answer:** B

**Explanation:**

References:

<https://cloudblogs.microsoft.com/sqlserver/2018/12/17/confidential-computing-using-always-encrypted-withsec>

**NEW QUESTION 7**

- (Exam Topic 3)

You need to optimize storage for CONT\_SQL3. What should you recommend?

- A. AlwaysOn
- B. Transactional processing
- C. General

D. Data warehousing

**Answer:** B

**Explanation:**

CONT\_SQL3 with the SQL Server role, 100 GB database size, Hyper-VM to be migrated to Azure VM. The storage should be configured to optimized storage for database OLTP workloads.

Azure SQL Database provides three basic in-memory based capabilities (built into the underlying database engine) that can contribute in a meaningful way to performance improvements:

In-Memory Online Transactional Processing (OLTP)

Clustered columnstore indexes intended primarily for Online Analytical Processing (OLAP) workloads Nonclustered columnstore indexes geared towards Hybrid Transactional/Analytical Processing (HTAP) workloads

References:

<https://www.databasejournal.com/features/mssql/overview-of-in-memory-technologies-of-azure-sqldatabase.htm>

**NEW QUESTION 8**

- (Exam Topic 3)

You are designing an Azure SQL Data Warehouse for a financial services company. Azure Active Directory will be used to authenticate the users.

You need to ensure that the following security requirements are met:

- ▶ Department managers must be able to create new database.
- ▶ The IT department must assign users to databases.
- ▶ Permissions granted must be minimized.

Which role memberships should you recommend? To answer, drag the appropriate roles to the correct groups. Each role 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.

Roles	Group	Role
dbmanager	Department managers	
loginmanager		
dc_admin	IT	
db_securityadmin		
db_owner		
db_accessadmin		

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: dbmanager

Members of the dbmanager role can create new databases. Box 2: db\_accessadmin

Members of the db\_accessadmin fixed database role can add or remove access to the database for Windows logins, Windows groups, and SQL Server logins.

References:

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

**NEW QUESTION 9**

- (Exam Topic 4)

You are designing an Azure SQL Data Warehouse. You plan to load millions of rows of data into the data warehouse each day.

You must ensure that staging tables are optimized for data loading. You need to design the staging tables.

What type of tables should you recommend?

- A. Round-robin distributed table
- B. Hash-distributed table
- C. Replicated table
- D. External table

**Answer:** A

**Explanation:**

To achieve the fastest loading speed for moving data into a data warehouse table, load data into a staging table. Define the staging table as a heap and use round-robin for the distribution option.

References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data>

**NEW QUESTION 10**

- (Exam Topic 4)

You are designing an Azure Databricks cluster that runs user-defined local processes. You need to recommend a cluster configuration that meets the following requirements:

- Minimize query latency.

- Reduce overall costs.
- Maximize the number of users that can run queries on the cluster at the same time. Which cluster type should you recommend?

- A. Standard with Autoscaling
- B. High Concurrency with Auto Termination
- C. High Concurrency with Autoscaling
- D. Standard with Auto Termination

**Answer:** C

**Explanation:**

High Concurrency clusters allow multiple users to run queries on the cluster at the same time, while minimizing query latency. Autoscaling clusters can reduce overall costs compared to a statically-sized cluster.

References:

<https://docs.azuredatabricks.net/user-guide/clusters/create.html> <https://docs.azuredatabricks.net/user-guide/clusters/high-concurrency.html#high-concurrency>  
<https://docs.azuredatabricks.net/user-guide/clusters/terminate.html> <https://docs.azuredatabricks.net/user-guide/clusters/sizing.html#enable-and-configure-autoscaling>

**NEW QUESTION 10**

- (Exam Topic 4)

A company has an application that uses Azure SQL Database as the data store.

The application experiences a large increase in activity during the last month of each year.

You need to manually scale the Azure SQL Database instance to account for the increase in data write operations.

Which scaling method should you recommend?

- A. Scale up by using elastic pools to distribute resources.
- B. Scale out by sharding the data across databases.
- C. Scale up by increasing the database throughput units.

**Answer:** C

**Explanation:**

As of now, the cost of running an Azure SQL database instance is based on the number of Database Throughput Units (DTUs) allocated for the database. When determining the number of units to allocate for the

solution, a major contributing factor is to identify what processing power is needed to handle the volume of expected requests.

Running the statement to upgrade/downgrade your database takes a matter of seconds.

**NEW QUESTION 11**

- (Exam Topic 4)

You are designing an application. You plan to use Azure SQL Database to support the application.

The application will extract data from the Azure SQL Database and create text documents. The text documents will be placed into a cloud-based storage solution.

The text storage solution must be accessible from an SMB network share.

You need to recommend a data storage solution for the text documents. Which Azure data storage type should you recommend?

- A. Queue
- B. Files
- C. Blob
- D. Table

**Answer:** B

**Explanation:**

Azure Files enables you to set up highly available network file shares that can be accessed by using the standard Server Message Block (SMB) protocol.

References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-introduction> <https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-overview>

**NEW QUESTION 15**

- (Exam Topic 4)

A company is developing a mission-critical line of business app that uses Azure SQL Database Managed Instance. You must design a disaster recovery strategy for the solution.

You need to ensure that the database automatically recovers when full or partial loss of the Azure SQL Database service occurs in the primary region.

What should you recommend?

- A. Failover-group
- B. Azure SQL Data Sync
- C. SQL Replication
- D. Active geo-replication

**Answer:** A

**Explanation:**

Auto-failover groups is a SQL Database feature that allows you to manage replication and failover of a group of databases on a SQL Database server or all databases in a Managed Instance to another region (currently in public preview for Managed Instance). It uses the same underlying technology as active geo-replication. You can initiate failover manually or you can delegate it to the SQL Database service based on a user-defined policy.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-auto-failover-group>

**NEW QUESTION 17**

- (Exam Topic 4)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You are designing an Azure SQL Database that will use elastic pools. You plan to store data about customers in a table. Each record uses a value for CustomerID. You need to recommend a strategy to partition data based on values in CustomerID. Proposed Solution: Separate data into customer regions by using horizontal partitioning. Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

We should use Horizontal Partitioning through Sharding, not divide through regions.

Note: Horizontal Partitioning - Sharding: Data is partitioned horizontally to distribute rows across a scaled out data tier. With this approach, the schema is identical on all participating databases. This approach is also called "sharding". Sharding can be performed and managed using (1) the elastic database tools libraries or (2)

self-sharding. An elastic query is used to query or compile reports across many shards.

References:

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

**NEW QUESTION 19**

- (Exam Topic 4)

You are developing a solution that performs real-time analysis of IoT data in the cloud. The solution must remain available during Azure service updates. You need to recommend a solution.

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

NOTE: Each correct selection is worth one point.

- A. Deploy an Azure Stream Analytics job to two separate regions that are not in a pair.
- B. Deploy an Azure Stream Analytics job to each region in a paired region.
- C. Monitor jobs in both regions for failure.
- D. Monitor jobs in the primary region for failure.
- E. Deploy an Azure Stream Analytics job to one region in a paired region.

**Answer: BC**

**Explanation:**

Stream Analytics guarantees jobs in paired regions are updated in separate batches. As a result there is a sufficient time gap between the updates to identify potential breaking bugs and remediate them.

Customers are advised to deploy identical jobs to both paired regions.

In addition to Stream Analytics internal monitoring capabilities, customers are also advised to monitor the jobs as if both are production jobs. If a break is identified to be a result of the Stream Analytics service update, escalate appropriately and fail over any downstream consumers to the healthy job output. Escalation to support will prevent the paired region from being affected by the new deployment and maintain the integrity of the paired jobs.

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-job-reliability>

**NEW QUESTION 22**

- (Exam Topic 4)

You are designing a Spark job that performs batch processing of daily web log traffic.

When you deploy the job in the production environment, it must meet the following requirements:

- Run once a day.
- Display status information on the company intranet as the job runs. You need to recommend technologies for triggering and monitoring jobs.

Which technologies should you recommend? To answer, drag the appropriate technologies to the correct locations. Each technology 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.

Technologies	Requirement	Technology
Livy	Triggering of jobs	<input type="text"/>
Beeline	Monitoring of jobs	<input type="text"/>
Azure Logic App		
Azure API App		

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Livy

You can use Livy to run interactive Spark shells or submit batch jobs to be run on Spark. Box 2: Beeline

Apache Beeline can be used to run Apache Hive queries on HDInsight. You can use Beeline with Apache Spark.

Note: Beeline is a Hive client that is included on the head nodes of your HDInsight cluster. Beeline uses JDBC to connect to HiveServer2, a service hosted on your

HDInsight cluster. You can also use Beeline to access Hive on HDInsight remotely over the internet.

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-livy-rest-interface> <https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/apache-hadoop-use-hive-beeline>

**NEW QUESTION 26**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your DP-201 Exam with Our Prep Materials Via below:**

<https://www.certleader.com/DP-201-dumps.html>