

➤ **Vendor:** Microsoft➤ **Exam Code:** DP-201➤ **Exam Name:** Implementing an Azure Data Solution➤ **New Updated Questions from** [Braindump2go](https://www.braindump2go.com) (Updated in [May/2020](#))**Visit Braindump2go and Download Full Version DP-201 Exam Dumps****QUESTION 115**

You need to design a telemetry data solution that supports the analysis of log files in real time.

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

NOTE: Each correct selection is worth one point.

- A. Azure Databricks
- B. Azure Data Factory
- C. Azure Event Hubs
- D. Azure Data Lake Storage Gen 2
- E. Azure IoT Hub

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You connect a data ingestion system with Azure Databricks to stream data into an Apache Spark cluster in near real-time. You set up data ingestion system using Azure Event Hubs and then connect it to Azure Databricks to process the messages coming through.


Note: Azure Event Hubs is a highly scalable data streaming platform and event ingestion service, capable of receiving and processing millions of events per second. Event Hubs can process and store events, data, or telemetry produced by distributed software and devices. Data sent to an event hub can be transformed and stored using any real-time analytics provider or batching/storage adapters.

References:

<https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-stream-from-eventhubs>

QUESTION 116

gn pattern based on the Lambda architecture as shown in the exhibit.



Which Azure service should you use for the hot path?

- A. Azure Databricks
- B. Azure SQL Database
- C. Azure Data Factory
- D. Azure Database for PostgreSQL

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In Azure, all of the following data stores will meet the core requirements supporting real-time processing:

Apache Spark in Azure Databricks

Azure Stream Analytics

HDInsight with Spark

Streaming HDInsight

with Storm

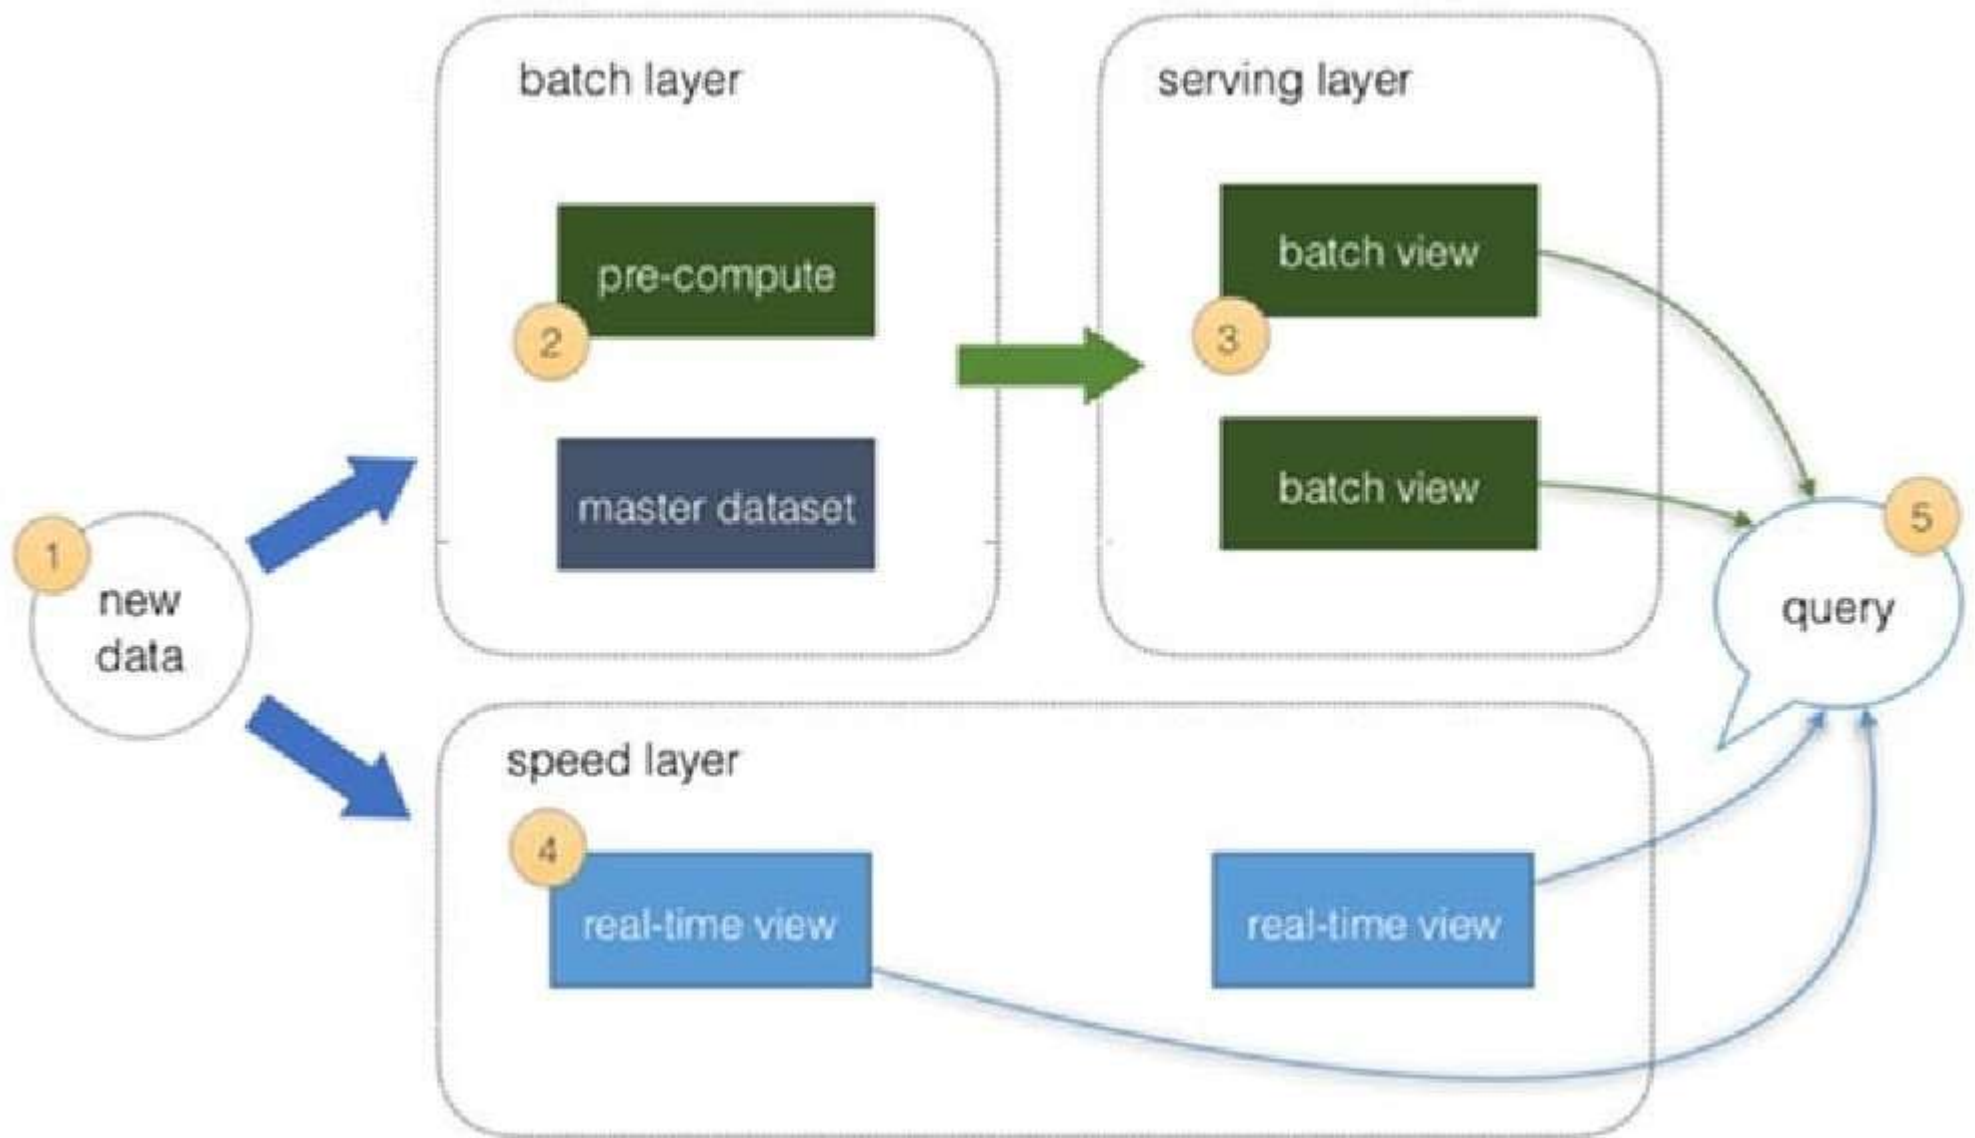
Azure Functions

Azure App Service WebJobs

Note: Lambda architectures use batch-processing, stream-processing, and a serving layer to minimize the latency involved in querying big data.

[DP-201 Exam Dumps](#) [DP-201 Exam Questions](#) [DP-201 PDF Dumps](#) [DP-201 VCE Dumps](#)

<https://www.braindump2go.com/dp-201.html>



References:

<https://azure.microsoft.com/en-us/blog/lambda-architecture-using-azure-cosmosdb-faster-performance-low-tco-low-devops/> <https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/stream-processing>

QUESTION 117

You are designing an audit strategy for an Azure SQL Database environment.

You need to recommend a solution to provide real-time notifications for potential security breaches. The solution must minimize development effort. Which destination should you include in the recommendation?

- A. Azure Blob storage
- B. Azure SQL Data Warehouse
- C. Azure Event Hubs
- D. Azure Log Analytics

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Auditing for Azure SQL Database and SQL Data Warehouse tracks database events and writes them to an audit log in your Azure storage account, Log Analytics workspace or Event Hubs.

Alerts in Azure Monitor can identify important information in your Log Analytics repository. They are created by alert rules that automatically run log searches at regular intervals, and if results of the log search match particular criteria, then an alert record is created and it can be configured to perform an automated response.

References:

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

<https://docs.microsoft.com/en-us/azure/azure-monitor/learn/tutorial-response>

QUESTION 118

You need to design a real-time stream solution that uses Azure Functions to process data uploaded to Azure Blob Storage.

The solution must meet the following requirements:

- Support up to 1 million blobs.
- Scaling must occur automatically.
- Costs must be minimized.

What should you recommend?

- A. Deploy the Azure Function in an App Service plan and use a Blob trigger.
- B. Deploy the Azure Function in a Consumption plan and use an Event Grid trigger.
- C. Deploy the Azure Function in a Consumption plan and use a Blob trigger.
- D. Deploy the Azure Function in an App Service plan and use an Event Grid trigger.

Correct Answer: C

[DP-201 Exam Dumps](#) [DP-201 Exam Questions](#) [DP-201 PDF Dumps](#) [DP-201 VCE Dumps](#)

<https://www.braindump2go.com/dp-201.html>

Section: (none)**Explanation****Explanation/Reference:****Explanation:**

Create a function, with the help of a blob trigger template, which is triggered when files are uploaded to or updated in Azure Blob storage.

You use a consumption plan, which is a hosting plan that defines how resources are allocated to your function app. In the default Consumption Plan, resources are added dynamically as required by your functions. In this serverless hosting, you only pay for the time your functions run. When you run in an App Service plan, you must manage the scaling of your function app.

References:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function>

QUESTION 119

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 SQL database that has columns. The columns contain sensitive Personally Identifiable Information (PII) data.

You need to design a solution that tracks and stores all the queries executed against the PII data. You must be able to review the data in Azure Monitor, and the data must be available for at least 45 days.

Solution: You create a SELECT trigger on the table in SQL Database that writes the query to a new table in the database, and then executes a stored procedure that looks up the column classifications and joins to the query text.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B**Section: (none)****Explanation****Explanation/Reference:****Explanation:**

Instead add classifications to the columns that contain sensitive data and turn on Auditing.

Note: Auditing has been enhanced to log sensitivity classifications or labels of the actual data that were returned by the query. This would enable you to gain insights on who is accessing sensitive data.

References:

<https://azure.microsoft.com/en-us/blog/announcing-public-preview-of-data-discovery-classification-for-microsoft-azure-sql-data-warehouse/>

QUESTION 120

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 SQL database that has columns. The columns contain sensitive Personally Identifiable Information (PII) data.

You need to design a solution that tracks and stores all the queries executed against the PII data. You must be able to review the data in Azure Monitor, and the data must be available for at least 45 days.

Solution: You add classifications to the columns that contain sensitive data. You turn on Auditing and set the audit log destination to use Azure Blob storage.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A**Section: (none)****Explanation****Explanation/Reference:****Explanation:**

Auditing has been enhanced to log sensitivity classifications or labels of the actual data that were returned by the query. This would enable you to gain insights on who is accessing sensitive data.

References:

<https://azure.microsoft.com/en-us/blog/announcing-public-preview-of-data-discovery-classification-for-microsoft-azure-sql-data-warehouse/>

QUESTION 121

You need to recommend a security solution for containers in Azure Blob storage. The solution must ensure that only read permissions are granted to a specific user for a specific container.

What should you include in the recommendation?

A. shared access signatures (SAS)

B. an RBAC role in Azure Active Directory (Azure AD)

C. public read access for blobs only

D. access keys

Correct Answer: A**Section: (none)****Explanation****Explanation/Reference:****Explanation:**

You can delegate access to read, write, and delete operations on blob containers, tables, queues, and file shares that are not permitted with a service SAS. Note: A shared access signature (SAS) provides secure delegated access to resources in your storage account without compromising the security of your data.

With a SAS, you have granular control over how a client can access your data. You can control what resources the client may access, what permissions they have on those resources, and how long the SAS is valid, among other parameters.

Incorrect Answers:

C: You can enable anonymous, public read access to a container and its blobs in Azure Blob storage. By doing so, you can grant read-only access to these resources without sharing your account key, and without requiring a shared access signature (SAS).

Public read access is best for scenarios where you want certain blobs to always be available for anonymous read access.

References:

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

QUESTION 122

You are designing the security for an Azure SQL database.

You have an Azure Active Directory (Azure AD) group named Group1.

You need to recommend a solution to provide Group1 with read access to the database only.

What should you include in the recommendation?

- A. a contained database user
- B. a SQL login
- C. an RBAC role
- D. a shared access signature (SAS)

Correct Answer: A Section:

(none) Explanation

Explanation/Reference:

Explanation:

Create a User for a security group

A best practice for managing your database is to use Windows security groups to manage user access.

That way you can simply manage the customer at the Security Group level in Active Directory granting appropriate permissions. To add a security group to SQL Data Warehouse, you use the Display Name of the security group as the principal in the CREATE USER statement.

CREATE USER [<Security Group Display Name>] FROM EXTERNAL PROVIDER WITH DEFAULT_SCHEMA = [<schema>];

In our AD instance, we have a security group called Sales Team with an alias of salesteam@company.com. To add this security group to SQL Data Warehouse you simply run the following statement:

CREATE USER [Sales Team] FROM EXTERNAL PROVIDER WITH DEFAULT_SCHEMA = [sales];

References:

https://blogs.msdn.microsoft.com/sql_dw/2017/07/28/adding-ad-users-and-security-groups-to-azure-sql-data-warehouse/

QUESTION 123

You store data in an Azure SQL data warehouse.

You need to design a solution to ensure that the data warehouse and the most current data is available within one hour of a datacenter failure.

Which three actions should you include in the design? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Each day, restore the data warehouse from a geo-redundant backup to an available Azure region.
- B. If a failure occurs, update the connection strings to point to the recovered data warehouse.
- C. If a failure occurs, modify the Azure Firewall rules of the data warehouse.
- D. Each day, create Azure Firewall rules that allow access to the restored data warehouse.
- E. Each day, restore the data warehouse from a user-defined restore point to an available Azure region.

Correct Answer: BDE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

E: You can create a user-defined restore point and restore from the newly created restore point to a new data warehouse in a different region.

Note: A data warehouse snapshot creates a restore point you can leverage to recover or copy your data warehouse to a previous state.

A data warehouse restore is a new data warehouse that is created from a restore point of an existing or deleted data warehouse. On average within the same region, restore rates typically take around 20 minutes.

Incorrect Answers:

A: SQL Data Warehouse performs a geo-backup once per day to a paired data center. The RPO for a geo-restore is 24 hours. You can restore the geo-backup to a server in any other region where SQL Data Warehouse is supported. A geo-backup ensures you can restore data warehouse in case you cannot access the restore points in your primary region.

References:

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

QUESTION 124

You are planning a big data solution in Azure.

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

- Be optimized for batch processing.
- Support autoscaling.
- Support per-cluster scaling.

Which technology should you recommend?

- A. Azure Data Warehouse
- B. Azure HDInsight with Spark
- C. Azure Analysis Services
- D. Azure Databricks

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Azure Databricks is an Apache Spark-based analytics platform. Azure Databricks supports autoscaling and manages the Spark cluster for you.

Incorrect Answers:

A, B:

Capability	Azure Data Lake Analytics	SQL Data Warehouse	HDInsight with Spark	HDInsight with Hive	HDInsight with Hive LLAP
Autoscaling	No	No	No	No	No

QUESTION 125

You are designing an Azure SQL data warehouse 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. column-level security
- D. Always Encrypted

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

SQL Database dynamic data masking limits sensitive data exposure by masking it to non-privileged users.

The Credit card masking method exposes the last four digits of the designated fields and adds a constant string as a prefix in the form of a credit card. Example:

XXXX-XXXX-XXXX-1234

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started>