Vendor: Microsoft

Exam Code: DP-200

Exam Name: Implementing an Azure Data Solution

Version: DEMO

QUESTION 1
Case Study 1 - Proseware, Inc

Background
Proseware, Inc, develops and manages a product named Poll Taker. The product is used for delivering public opinion polling and analysis.

You need to ensure that phone-based poling data can be analyzed in the PollingData database.

How should you configure Azure Data Factory?

A. Use a tumbling schedule trigger
B. Use an event-based trigger
C. Use a schedule trigger
D. Use manual execution

Answer: C
Explanation:
When creating a schedule trigger, you specify a schedule (start date, recurrence, end date etc.) for the trigger, and associate with a Data Factory pipeline.
Scenario:
All data migration processes must use Azure Data Factory
All data migrations must run automatically during non-business hours
References:
https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-schedule-trigger

QUESTION 2
Case Study 2 - Contoso

Overview
Current environment
Contoso relies on an extensive partner network for marketing, sales, and distribution. Contoso uses external companies that manufacture everything from the actual pharmaceutical to the packaging.

The majority of the company's data reside in Microsoft SQL Server database. Application databases fall into one of the following tiers:

Applications	Tier	Replication	Notes
Internal Contoso	1	Yes	
Internal Contoso	2	SQL Data Sync	Data Sync between databases
Internal Partner	3	Yes	Replicate to Partner
External Contoso	4,5,6	Yes	
External Partner	7,8	No	Partner managed
Internal Distribution and Sales	9	Yes, once ingested at branches	Data ingested from Contoso branches
External Distribution and Sales	10	Yes, once ingested at Contoso main office	Data is ingested from multiple sources

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 questions 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 need to implement diagnostic logging for Data Warehouse monitoring.

Which log should you use?

A. RequestSteps
B. DmsWorkers
C. SqlRequests
D. ExecRequests

Answer: C
Explanation:
Scenario:
The Azure SQL Data Warehouse cache must be monitored when the database is being used.

Metric	Description
A	Low cache hit %, high cache usage %
B	Low cache hit %, low cache usage %
C	High cache hit %, high cache usage %

References:
https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-sql-requests-transact-sq

QUESTION 3
You are a data engineer implementing a lambda architecture on Microsoft Azure.
You use an open-source big data solution to collect, process, and maintain data.
The analytical data store performs poorly.

You must implement a solution that meets the following requirements:

```
- Provide data warehousing
- Reduce ongoing management activities
- Deliver SQL query responses in less than one second
```

You need to create an HDInsight cluster to meet the requirements.

Which type of cluster should you create?

A. Interactive Query
B. Apache Hadoop
C. Apache HBase
D. Apache Spark

Answer: D
Explanation:
Lambda Architecture with Azure:
Azure offers you a combination of following technologies to accelerate real-time big data analytics:
1. Azure Cosmos DB, a globally distributed and multi-model database service.

---

2. Apache Spark for Azure HDInsight, a processing framework that runs large-scale data analytics applications.
3. Azure Cosmos DB change feed, which streams new data to the batch layer for HDInsight to process.
4. The Spark to Azure Cosmos DB Connector
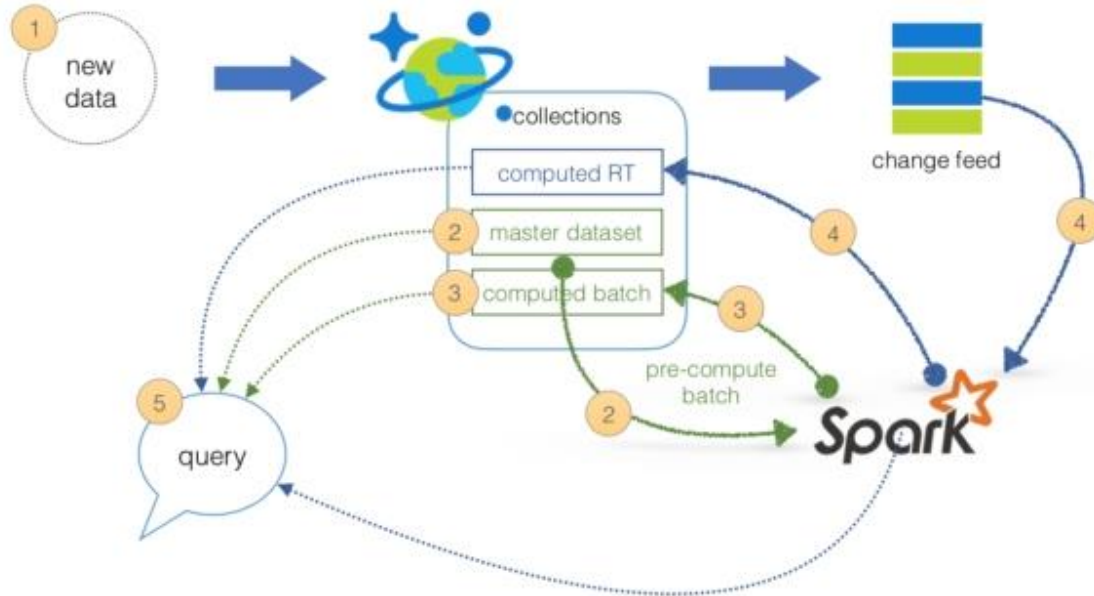
Note: Lambda architecture is a data-processing architecture designed to handle massive quantities of data by taking advantage of both batch processing and stream processing methods, and minimizing the latency involved in querying big data.
References:
https://sqlwithmanoj.com/2018/02/16/what-is-lambda-architecture-and-what-azure-offers-with-its-new-cosmos-db/

QUESTION 4
You develop data engineering solutions for a company. The company has on-premises Microsoft SQL Server databases at multiple locations.

The company must integrate data with Microsoft Power BI and Microsoft Azure Logic Apps. The solution must avoid single points of failure during connection and transfer to the cloud. The solution must also minimize latency.

You need to secure the transfer of data between on-premises databases and Microsoft Azure.

What should you do?

A. Install a standalone on-premises Azure data gateway at each location
B. Install an on-premises data gateway in personal mode at each location
C. Install an Azure on-premises data gateway at the primary location
D. Install an Azure on-premises data gateway as a cluster at each location

Answer: D
Explanation:

You can create high availability clusters of On-premises data gateway installations, to ensure your organization can access on-premises data resources used in Power BI reports and dashboards. Such clusters allow gateway administrators to group gateways to avoid single points of failure in accessing on-premises data resources. The Power BI service always uses the primary gateway in the cluster, unless it's not available. In that case, the service switches to the next gateway in the cluster, and so on.
References:
https://docs.microsoft.com/en-us/power-bi/service-gateway-high-availability-clusters

QUESTION 5
You are a data architect. The data engineering team needs to configure a synchronization of data between an on-premises Microsoft SQL Server database to Azure SQL Database.

Ad-hoc and reporting queries are being overutilized the on-premises production instance. The synchronization process must:

```
- Perform an initial data synchronization to Azure SQL Database with
minimal downtime
- Perform bi-directional data synchronization after initial
synchronization
```

You need to implement this synchronization solution.

Which synchronization method should you use?

A. transactional replication
B. Data Migration Assistant (DMA)
C. backup and restore
D. SQL Server Agent job
E. Azure SQL Data Sync

Answer: E
Explanation:
SQL Data Sync is a service built on Azure SQL Database that lets you synchronize the data you select bi-directionally across multiple SQL databases and SQL Server instances.
With Data Sync, you can keep data synchronized between your on-premises databases and Azure SQL databases to enable hybrid applications.
Compare Data Sync with Transactional Replication

	Data Sync	Transactional Replication
Advantages	- Active-active support - Bi-directional between on-premises and Azure SQL Database	- Lower latency - Transactional consistency - Reuse existing topology after migration
Disadvantages	- 5 min or more latency - No transactional consistency - Higher performance impact	- Can't publish from Azure SQL Database single database or pooled database - High maintenance cost

References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-sync-data

Thank You for Trying Our Product

PassLeader Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.passleader.com/all-products.html

10% Discount Coupon Code: STNAR2014