Microsoft Data Certification DP-201 Exam

- **Vendor: Microsoft**
- **Exam Code: DP-201**
- **Exam Name: Designing an Azure Data Solution**

NEW QUESTION 1
You are evaluating data storage solutions to support a new application. You need to recommend a data storage solution that represents data by using nodes and relationships in graph structures. Which data storage solution should you recommend?

A. Blob Storage
B. Cosmos DB
C. Data Lake Store
D. HDInsight

Answer: B
Explanation:
For large graphs with lots of entities and relationships, you can perform very complex analyses very quickly. Many graph databases provide a query language that you can use to traverse a network of relationships efficiently.
https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/data-store-overview

NEW QUESTION 2
You are designing a data processing solution that will implement the lambda architecture pattern. The solution will use Spark running on HDInsight for data processing. You need to recommend a data storage technology for the solution. Which two technologies should you recommend? (Each correct answer presents a complete solution. Choose two.)

A. Azure Cosmos DB
B. Azure Service Bus
C. Azure Storage Queue
D. Apache Cassandra
E. Kafka HDInsight

Answer: AE
Explanation:
Option A:
To implement a lambda architecture on Azure, you can combine the following technologies to accelerate real-time big data analytics:
- Azure Cosmos DB, the industry's first globally distributed, multi-model database service.
- Apache Spark for Azure HDInsight, a processing framework that runs large-scale data analytics applications.
- Azure Cosmos DB change feed, which streams new data to the batch layer for HDInsight to process.
- The Spark to Azure Cosmos DB Connector.
Option E: You can use Apache Spark to stream data into or out of Apache Kafka on HDInsight using DStreams.
https://docs.microsoft.com/en-us/azure/cosmos-db/lambda-architecture

NEW QUESTION 3
You are designing a solution for a company. The solution will use model training for objective classification. You need to design the solution. What should you recommend?

A. An Azure Cognitive Services application.
B. A Spark Streaming job.
C. Interactive Spark queries.
D. Power BI models.
E. A Spark application that uses Spark MLlib.

Answer: E
Explanation:
Spark in SQL Server big data cluster enables AI and machine learning. You can use Apache Spark MLlib to create a machine learning application to do simple predictive analysis on an open dataset. MLlib is a core Spark library that provides many utilities useful for machine learning tasks, including utilities that are suitable for:
- Classification
- Regression
- Clustering
- Topic modeling
- Singular value decomposition (SVD) and principal component analysis (PCA)
- Hypothesis testing and calculating sample statistics
https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-machine-learning-mllib-ipython

NEW QUESTION 4
A company is developing a solution to manage inventory data for a group of automotive repair shops. The solution will use Azure SQL Data Warehouse as the data store. Shops will upload data every 10 days. Data corruption checks must run each time data is uploaded. If corruption is detected, the corrupted data must be removed. You need to ensure that upload processes and data corruption checks do not impact reporting and analytics processes that use the data warehouse.
Solution: Insert data from shops and perform the data corruption check in a transaction. Rollback transfer if corruption is detected.
Does the solution meet the goal?

A. Yes
B. No

Answer: B
Explanation:
Instead, create a user-defined restore point before data is uploaded. Delete the restore point after data corruption checks complete.
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore

NEW QUESTION 5
A company is developing a solution to manage inventory data for a group of automotive repair shops. The solution will use Azure SQL Data Warehouse as the data store. Shops will upload data every 10 days. Data corruption checks must run each time data is uploaded. If corruption is detected, the corrupted data must be removed. You need to ensure that upload processes and data corruption checks do not impact reporting and analytics processes that use the data warehouse.
Solution: Create a user-defined restore point before data is uploaded. Delete the restore point after data corruption checks complete.
Does the solution meet the goal?

A. Yes
B. No

Answer: A
Explanation:
User-Defined Restore Points. This feature enables you to manually trigger snapshots to create restore points of your data warehouse before and after large modifications. This capability ensures that restore points are logically consistent, which provides additional data protection in case of any workload interruptions or user errors for quick recovery time.

Note: A data warehouse restore is a new data warehouse that is created from a restore point of an existing or deleted data warehouse. Restoring your data warehouse is an essential part of any business continuity and disaster recovery strategy because it re-creates your data after accidental corruption or deletion.
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore

NEW QUESTION 6
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.
Incorrect:
Not B: Consider that loading is usually a two-step process in which you first load to a staging table and then insert the data into a production data warehouse table. If the production table uses a hash distribution, the total time to load and insert might be faster if you define the staging table with the hash distribution. Loading to the staging table takes longer, but the second step of inserting the rows to the production table does not incur data movement across the distributions.
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data

NEW QUESTION 7
Drag and Drop
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.)

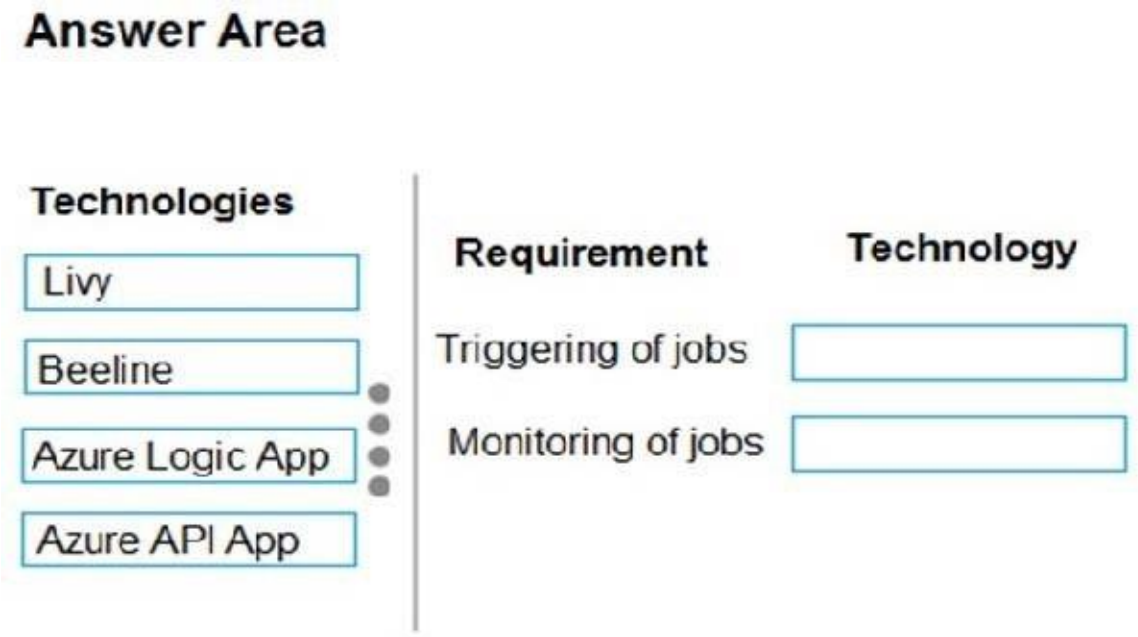

Answer:

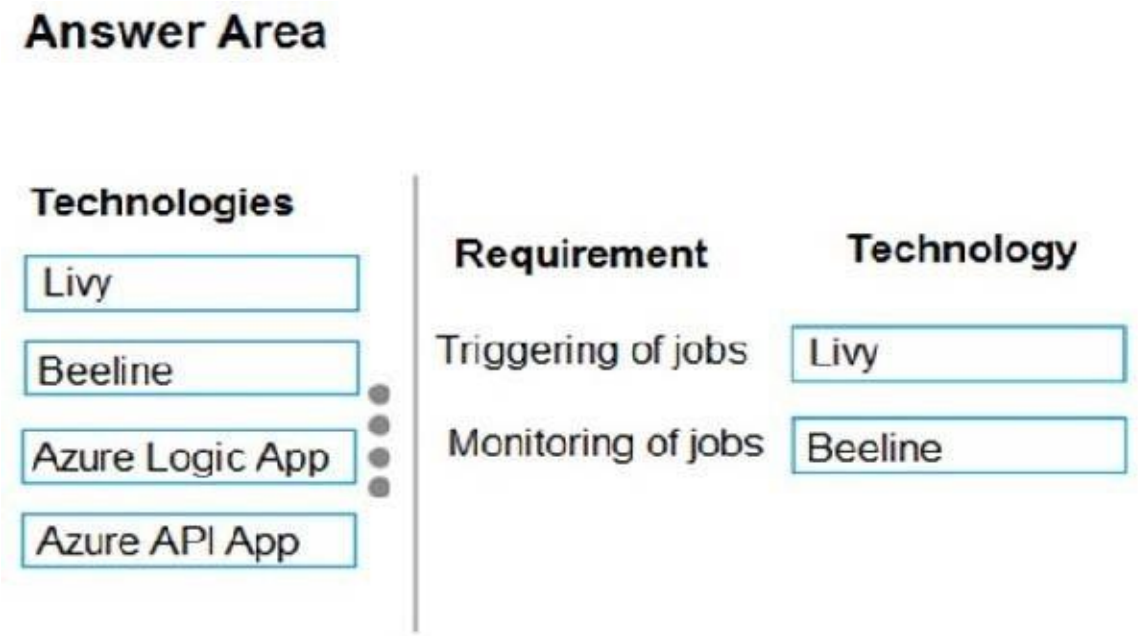

Explanation:
Box 1: Livy. You can use Livy to run interactive Spark shells or submit batch jobs to be run on Spark.
Box 2: Beeline. Apache Beeline can be used to run Apache Hive queries on HDInsight. You can use Beeline with Apache Spark.
Note: Beeline is a Hive client that is included on the head nodes 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.
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 8
HotSpot
You are designing a recovery strategy for your Azure SQL Databases. The recovery strategy must use default automated backup settings. The solution must include a Point-in time restore recovery strategy. You need to recommend which backups to use and the order in which to restore backups.
What should you recommend? (To answer, select the appropriate configuration in the answer area.)

Answer Area

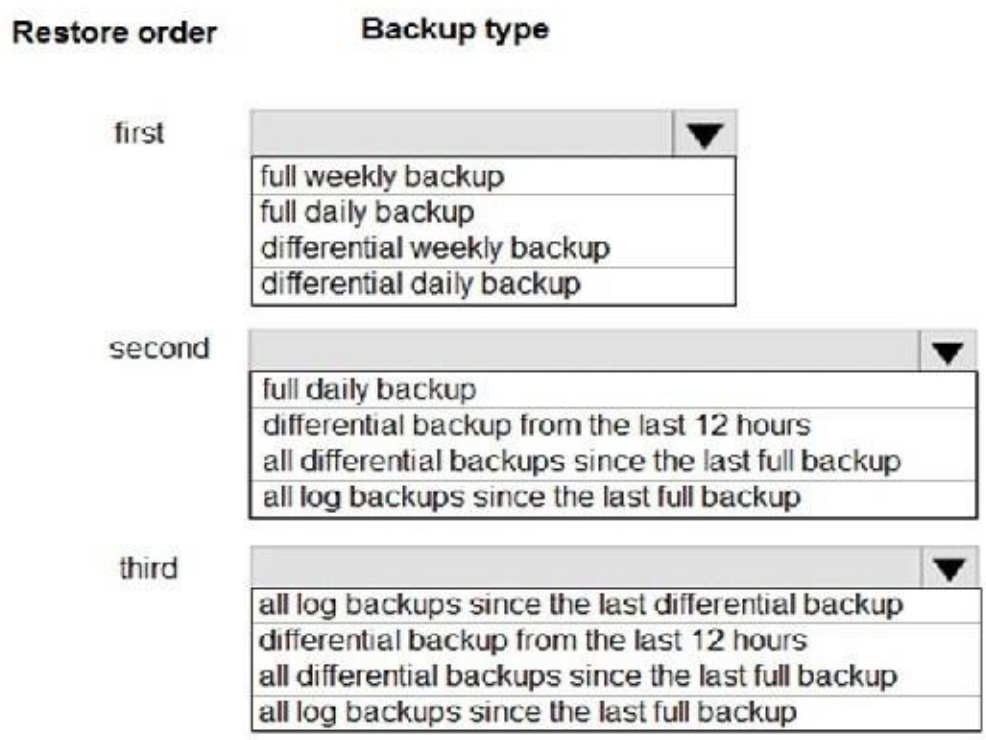

Answer:

Answer Area

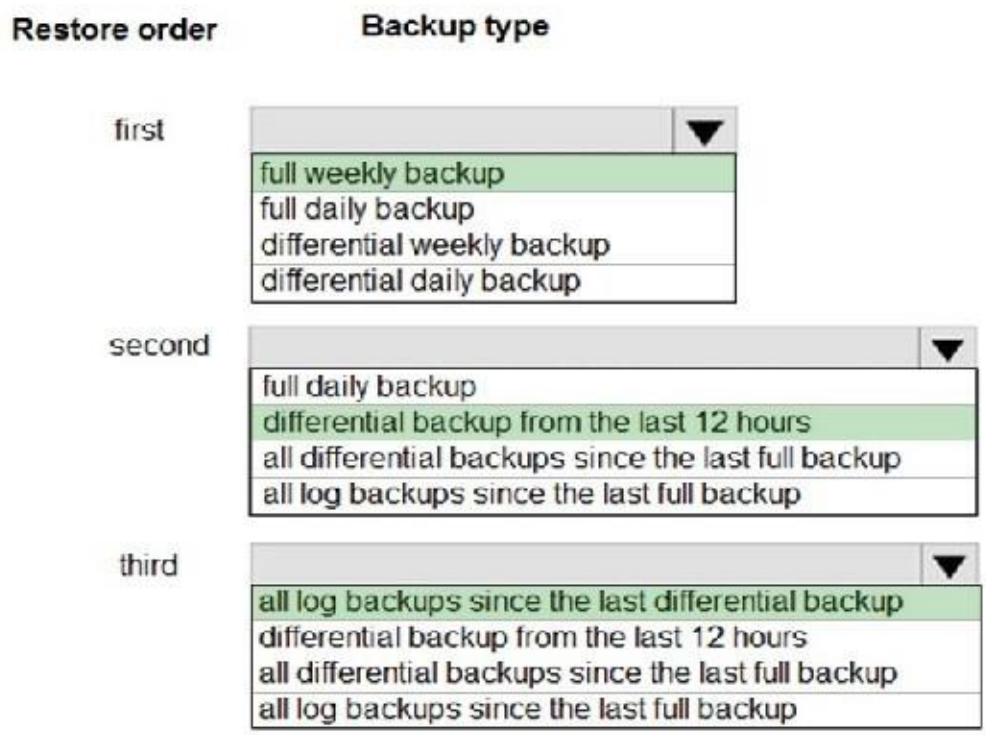

Explanation:
All Basic, Standard, and Premium databases are protected by automatic backups. Full backups are taken every week, differential backups every day, and log backups every 5 minutes.
https://azure.microsoft.com/sv-se/blog/azure-sql-database-point-in-time-restore/

NEW QUESTION 9
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:
Not A: Transparent Data Encryption (TDE) encrypts SQL Server, Azure SQL Database, and Azure SQL Data Warehouse data files, known as encrypting data at rest. TDE does not provide encryption across communication channels.
https://cloudblogs.microsoft.com/sqlserver/2018/12/17/confidential-computing-using-always-encrypted-with-secure-enclaves-in-sql-server-2019-preview/

NEW QUESTION 10
You plan to use Azure SQL Database to support a line of business app. You need to identify sensitive data that is stored in the database and monitor access to the data. Which three actions

should you recommend? (Each correct answer presents part of the solution. Choose three.)

A. Enable Data Discovery and Classification.
B. Implement Transparent Data Encryption (TDE).
C. Enable Auditing.
D. Run Vulnerability Assessment.
E. Use Advanced Threat Protection.

Answer: CDE

NEW QUESTION 11
.....

Case Study 1 - Trey Research
Trey Research is a technology innovator. The company partners with regional transportation department office to build solutions that improve traffic flow and safety. The company is developing the following solutions:
.....

NEW QUESTION 41
You need to design the vehicle images storage solution. What should you recommend?

A. Azure Media Services
B. Azure Premium Storage Account
C. Azure Redis Cache
D. Azure Cosmos DB

Answer: B
Explanation:
Premium Storage stores data on the latest technology Solid State Drives (SSDs) whereas Standard Storage stores data on Hard Disk Drives (HDDs). Premium Storage is designed for Azure Virtual Machine workloads which require consistent high IO performance and low latency in order to host IO intensive workloads like OLTP, Big Data, and Data Warehousing on platforms like SQL Server, MongoDB, Cassandra, and others. With Premium Storage, more customers will be able to lift-and-shift demanding enterprise applications to the cloud.
https://azure.microsoft.com/es-es/blog/introducing-premium-storage-high-performance-storage-for-azure-virtual-machine-workloads/

NEW QUESTION 42
You need to recommend an Azure SQL Database pricing tier for Planning Assistance. Which pricing tier should you recommend?

A. Business critical Azure SQL Database single database.
B. General purpose Azure SQL Database Managed Instance.
C. Business critical Azure SQL Database Managed Instance.
D. General purpose Azure SQL Database single database.

Answer: B
Explanation:
Azure resource costs must be minimized where possible. Data used for Planning Assistance must be stored in a sharded Azure SQL Database. The SLA for Planning Assistance is 70 percent, and multiday outages are permitted.

NEW QUESTION 43
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 44
.....

Case Study 2 - Proseware, Inc.
Proseware, Inc. identifies the following business requirements:
- You must transfer all images and customer data to cloud storage and remove on-premises servers.
- You must develop an analytical processing solution for transforming customer data.
- You must develop an image object and color tagging solution.
- Capital expenditures must be minimized.
- Cloud resource costs must be minimized.
The solution has the following technical requirements:
.....

NEW QUESTION 51
You need to design the solution for analyzing customer data. What should you recommend?

A. Azure Databricks
B. Azure Data Lake Storage
C. Azure SQL Data Warehouse
D. Azure Cognitive Services
E. Azure Batch

Answer: A
Explanation:
Customer data must be analyzed using managed Spark clusters. You create spark clusters through Azure Databricks.
https://docs.microsoft.com/en-us/azure/azure-databricks/quickstart-create-databricks-workspace-portal

NEW QUESTION 52
You need to recommend a solution for storing the image tagging data. What should you recommend?

A. Azure File Storage
B. Azure Cosmos DB
C. Azure Blob Storage
D. Azure SQL Database
E. Azure SQL Data Warehouse

Answer: C
Explanation:
Image data must be stored in a single data store at minimum cost.
Note: Azure Blob storage is Microsoft's object storage solution for the cloud. Blob storage is optimized for storing massive amounts of unstructured data. Unstructured data is data that does not adhere to a particular data model or definition, such as text or binary data. Blob storage is designed for:

- Serving images or documents directly to a browser.
- Storing files for distributed access.
- Streaming video and audio.
- Writing to log files.
- Storing data for backup and restore, disaster recovery, and archiving.
- Storing data for analysis by an on-premises or Azure-hosted service.

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blobs-introduction

NEW QUESTION 53
You need to design a backup solution for the processed customer data. What should you include in the design?

A. AzCopy
B. AdlCopy
C. Geo-Redundancy
D. Geo-Replication

Answer: C
Explanation:
Geo-redundant storage (GRS) is designed to provide at least 99.99999999999999% (16 9's) durability of objects over a given year by replicating your data to a secondary region that is hundreds of miles away from the primary region. If your storage account has GRS enabled, then your data is durable even in the case of a complete regional outage or a disaster in which the primary region isn't recoverable.
https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs

NEW QUESTION 54
.....

Case Study 3 - Contoso
Contoso has the following virtual machines (VMs):
.....

NEW QUESTION 61
You need to design a solution to meet the SQL Server storage requirements for CONT_SQL3. Which type of disk should you recommend?

A. Standard SSD Managed Disk
B. Premium SSD Managed Disk
C. Ultra SSD Managed Disk

Answer: C
Explanation:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-types

NEW QUESTION 62
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
https://www.databasejournal.com/features/mssql/overview-of-in-memory-technologies-of-azure-sql-database.html

NEW QUESTION 63
You need to recommend a backup strategy for CONT_SQL1 and CONT_SQL2. What should you recommend?

A. Use AzCopy and store the data in Azure.
B. Configure Azure SQL Database long-term retention for all databases.
C. Configure Accelerated Database Recovery.
D. Use DWLoader.

Answer: B

NEW QUESTION 64
.....

NEW QUESTION 111
You are designing an application that will have an Azure virtual machine. The virtual machine will access an Azure SQL database. The database will not be accessible from the Internet. You need to recommend a solution to provide the required level of access to the database. What should you include in the recommendation?

A. Deploy an On-premises data gateway.
B. Add a virtual network to the Azure SQL server that hosts the database.
C. Add an application gateway to the virtual network that contains the Azure virtual machine.
D. Add a virtual network gateway to the virtual network that contains the Azure virtual machine.

Answer: B
Explanation:
When you create an Azure virtual machine (VM), you must create a virtual network (VNet) or use an existing VNet. You also need to decide how your VMs are intended to be accessed on the VNet.
Incorrect:
Not C: Azure Application Gateway is a web traffic load balancer that enables you to manage traffic to your web applications.
Not D: A VPN gateway is a specific type of virtual network gateway that is used to send encrypted traffic between an Azure virtual network and an on-premises location over the public Internet.
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/network-overview

NEW QUESTION 112
You are designing a data store that will store organizational information for a company. The data will be used to identify the relationships between users. The data will be stored in an Azure Cosmos DB database and will contain several million objects. You need to recommend which API to use for the database. The API must minimize the complexity to query the user relationships. The solution must support fast traversals. Which API should you recommend?

A. MongoDB
B. Table
C. Gremlin
D. Cassandra

Answer: C
Explanation:
Gremlin features fast queries and traversals with the most widely adopted graph query standard.
https://docs.microsoft.com/th-th/azure/cosmos-db/graph-introduction?view=azurermps-5.7.0

NEW QUESTION 113
You need to recommend a storage solution to store flat files and columnar optimized files. The solution must meet the following requirements:
- Store standardized data that data scientists will explore in a curated folder.
- Ensure that applications cannot access the curated folder.
- Store staged data for import to applications in a raw folder.
- Provide data scientists with access to specific folders in the raw folder and all the content the curated folder.
Which storage solution should you recommend?

A. Azure SQL Data Warehouse
B. Azure Blob storage
C. Azure Data Lake Storage Gen2
D. Azure SQL Database

Answer: B
Explanation:
Azure Blob Storage containers is a general purpose object store for a wide variety of storage scenarios. Blobs are stored in containers, which are similar to folders.
Incorrect:
Not C: Azure Data Lake Storage is an optimized storage for big data analytics workloads.
https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/data-storage

NEW QUESTION 114
You have a MongoDB database that you plan to migrate to an Azure Cosmos DB account that uses the MongoDB API. During testing, you discover that the migration takes longer than expected. You need to recommend a solution that will reduce the amount of time it takes to migrate the data. What are two possible recommendations to achieve this goal? (Each correct answer presents a complete solution. Choose two.)

A. Increase the Request Units (RUs).
B. Turn off indexing.
C. Add a write region.
D. Create unique indexes.
E. Create compound indexes.

Answer: AB
Explanation:
A: Increase the throughput during the migration by increasing the Request Units (RUs). For customers that are migrating many collections within a database, it is strongly recommend to configure database-level throughput. You must make this choice when you create the database. The minimum database-level throughput capacity is 400 RU/sec. Each collection sharing database-level throughput requires at least 100 RU/sec.
B: By default, Azure Cosmos DB indexes all your data fields upon ingestion. You can modify the indexing policy in Azure Cosmos DB at any time. In fact, it is often recommended to turn off indexing

when migrating data, and then turn it back on when the data is already in Cosmos DB.
https://docs.microsoft.com/bs-latn-ba/Azure/cosmos-db/mongodb-pre-migration

NEW QUESTION 115
You need to recommend a storage solution for a sales system that will receive thousands of small files per minute. The files will be in JSON, text, and CSV formats. The files will be processed and transformed before they are loaded into an Azure data warehouse. The files must be stored and secured in folders. Which storage solution should you recommend?

A. Azure Data Lake Storage Gen2
B. Azure Cosmos DB
C. Azure SQL Database
D. Azure Blob storage

Answer: A
Explanation:
Azure provides several solutions for working with CSV and JSON files, depending on your needs. The primary landing place for these files is either Azure Storage or Azure Data Lake Store.1 Azure Data Lake Storage is an optimized storage for big data analytics workloads.
Incorrect:
Not D: Azure Blob Storage containers is a general purpose object store for a wide variety of storage scenarios. Blobs are stored in containers, which are similar to folders.
https://docs.microsoft.com/en-us/azure/architecture/data-guide/scenarios/csv-and-json

NEW QUESTION 116
You are designing an Azure Cosmos DB database that will support vertices and edges. Which Cosmos DB API should you include in the design?

A. SQL
B. Cassandra
C. Gremlin
D. Table

Answer: C
Explanation:
The Azure Cosmos DB Gremlin API can be used to store massive graphs with billions of vertices and edges.
https://docs.microsoft.com/en-us/azure/cosmos-db/graph-introduction

NEW QUESTION 117
You plan to store delimited text files in an Azure Data Lake Storage account that will be organized into department folders. You need to configure data access so that users see only the files in their respective department folder.
Solution: From the storage account, you enable a hierarchical namespace, and you use RBAC.
Does this meet the goal?

A. Yes
B. No

Answer: B
Explanation:
Disable the hierarchical namespace. And instead of RBAC use access control lists (ACLs). Note: Azure Data Lake Storage implements an access control model that derives from HDFS, which in turn derives from the POSIX access control model. Blob container ACLs does not support the hierarchical namespace, so it must be disabled.

https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-known-issues
https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-access-control

NEW QUESTION 118
You plan to store delimited text files in an Azure Data Lake Storage account that will be organized into department folders. You need to configure data access so that users see only the files in their respective department folder.
Solution: From the storage account, you disable a hierarchical namespace, and you use access control lists (ACLs).
Does this meet the goal?

A. Yes
B. No

Answer: A
Explanation:
Azure Data Lake Storage implements an access control model that derives from HDFS, which in turn derives from the POSIX access control model. Blob container ACLs does not support the hierarchical namespace, so it must be disabled.
https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-known-issues
https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-access-control

NEW QUESTION 119
You are designing a data storage solution for a database that is expected to grow to 50 TB. The usage pattern is singleton inserts, singleton updates, and reporting. Which storage solution should you use?

A. Azure SQL. Database elastic pools.
B. Azure SQL Data Warehouse.
C. Azure Cosmos DB that uses the Gremlin API.
D. Azure SQL Database Hyperscale.

Answer: D
Explanation:
A Hyperscale database is an Azure SQL database in the Hyperscale service tier that is backed by the Hyperscale scale-out storage technology. A Hyperscale database supports up to 100 TB of data and provides high throughput and performance, as well as rapid scaling to adapt to the workload requirements. Scaling is transparent to the application connectivity, query processing, etc. work like any other Azure SQL database.
Incorrect:
Not A: 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 Azure SQL Database server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.
Not B: Rather than SQL Data Warehouse, consider other options for operational (OLTP) workloads that have large numbers of singleton selects.
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tier-hyperscale-faq

NEW QUESTION 120
You are designing an Azure Databricks interactive cluster. The cluster will be used infrequently and will be configured for auto-termination. You need to ensure that the cluster configuration is retained indefinitely after the cluster is terminated. The solution must minimize costs. What should you do?

A. Clone the cluster after it is terminated.
B. Terminate the cluster manually when processing completes.
C. Create an Azure runbook that starts the cluster every 90 days.
D. Pin the cluster.

Answer: D
Explanation:
To keep an interactive cluster configuration even after it has been terminated for more than 30 days, an administrator can pin a cluster to the cluster list.
https://docs.azuredatabricks.net/clusters/clusters-manage.html#automatic-termination

NEW QUESTION 121
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

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

NEW QUESTION 122
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

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

NEW QUESTION 123
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

Answer: A
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:
Not 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.
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

NEW QUESTION 124
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. Choose three.)

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.

Answer: BDE
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:
Not 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.
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore

NEW QUESTION 125
You design data engineering solutions for a company that has locations around the world. You plan to deploy a large set of data to Azure Cosmos DB. The data must be accessible from all company locations. You need to recommend a strategy for deploying the data that minimizes latency for data read operations and minimizes costs. What should you recommend?

A. Use a single Azure Cosmos DB account. Enable multi-region writes.
B. Use a single Azure Cosmos DB account Configure data replication.
C. Use multiple Azure Cosmos DB accounts. For each account, configure the location to the closest Azure datacenter.
D. Use a single Azure Cosmos DB account. Enable geo-redundancy.
E. Use multiple Azure Cosmos DB accounts. Enable multi-region writes.

Answer: A
Explanation:
With Azure Cosmos DB, you can add or remove the regions associated with your account at any time. Multi-region accounts configured with multiple-write regions will be highly available for both writes and reads. Regional failovers are instantaneous and don't require any changes from the application.
https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability

NEW QUESTION 126
.....