

QUESTIONS & ANSWERS

Kill your exam at first Attempt



*BDS-C00 Dumps
BDS-C00 Braindumps
BDS-C00 Real Questions
BDS-C00 Practice Test
BDS-C00 dumps free*



Amazon

BDS-C00

AWS Certified Big Data Specialty (BDS-C00)



<http://killexams.com/pass4sure/exam-detail/BDS-C00>

Question: 75

An organization currently runs a large Hadoop environment in their data center and is in the process of creating an alternative Hadoop environment on AWS, using Amazon EMR. They generate around 20 TB of data on a monthly basis. Also on a monthly basis, files need to be grouped and copied to Amazon S3 to be used for the Amazon EMR environment. They have multiple S3 buckets across AWS accounts to which data needs to be copied. There is a 10G AWS Direct Connect setup between their data center and AWS, and the network team has agreed to allocate 50% of AWS Direct Connect bandwidth to data transfer. The data transfer cannot take more than two days. What would be the MOST efficient approach to transfer data to AWS on a monthly basis?

- A. Use an offline copy method, such as an AWS Snowball device, to copy and transfer data to Amazon S3.
- B. Configure a multipart upload for Amazon S3 on AWS Java SDK to transfer data over AWS Direct Connect.
- C. Use Amazon S3 transfer acceleration capability to transfer data to Amazon S3 over AWS Direct Connect.
- D. Setup S3DistCop tool on the on-premises Hadoop environment to transfer data to Amazon S3 over AWS Direct Connect.

Answer: B

Question: 76

An organization is developing a mobile social application and needs to collect logs from all devices on which it is installed. The organization is evaluating the Amazon Kinesis Data Streams to push logs and Amazon EMR to process data. They want to store data on HDFS using the default replication factor to replicate data among the cluster, but they are concerned about the durability of the data. Currently, they are producing 300 GB of raw data daily, with additional spikes during special events. They will need to scale out the Amazon EMR cluster to match the increase in streamed data. Which solution prevents data loss and matches compute demand?

- A. Use multiple Amazon EBS volumes on Amazon EMR to store processed data and scale out the Amazon EMR cluster as needed.
- B. C. Use Amazon DynamoD to store processed data and scale out the Amazon EMR cluster as needed.
- C. Use Amazon DynamoDB to store processed data and scale out the Amazon EMR cluster as needed.
- D. use Amazon Kinesis Data Firehose and, instead of using Amazon EMR, stream logs directly into Amazon Elasticsearch Service.

Answer: D

Question: 77

An advertising organization uses an application to process a stream of events that are received from clients in multiple unstructured formats. The application does the following: ? Transforms the events into a single structured format and streams them to Amazon Kinesis for real-time analysis. ? Stores the unstructured raw events from the log files on local hard drivers that are rotated and uploaded to Amazon S3. The organization wants to extract campaign performance reporting using an existing Amazon redshift cluster. Which solution will provide the

performance data with the LEAST number of operations?

- A. Install the Amazon Kinesis Data Firehose agent on the application servers and use it to stream the log files directly to Amazon Redshift.
- B. Create an external table in Amazon Redshift and point it to the S3 bucket where the unstructured raw events are stored.
- C. Write an AWS Lambda function that triggers every hour to load the new log files already in S3 to Amazon Redshift.
- D. Connect Amazon Kinesis Data Firehose to the existing Amazon Kinesis stream and use it to stream the event directly to Amazon Redshift.

Answer: B

Question: 78

An organization needs to store sensitive information on Amazon S3 and process it through Amazon EMR. Data must be encrypted on Amazon S3 and Amazon EMR at rest and in transit. Using Thrift Server, the Data Analysis team uses HIVE to interact with this data. The organization would like to grant access to only specific databases and tables, giving permission only to the SELECT statement. Which solution will protect the data and limit user access to the SELECT statement on a specific portion of data?

- A. Configure Transparent Data Encryption on Amazon EMR. Create an Amazon EC2 instance and install Apache Ranger. Configure the authorization on the cluster to use Apache Ranger.
- B. Configure data encryption at rest for EMR File System (EMRFS) on Amazon S3. Configure data encryption in transit for traffic between Amazon S3 and EMRFS. Configure storage and SQL base authorization on HiveServer2.
- C. D. Configure Security Group on Amazon EMR. Create an Amazon VPC endpoint for Amazon S3. Configure HiveServer2 to use Kerberos authentication on the cluster.
- D. Configure Security Group on Amazon EMR. Create an Amazon VPC endpoint for Amazon S3. Configure HiveServer2 to use Kerberos authentication on the cluster.

Answer: C

Question: 79

Multiple rows in an Amazon Redshift table were accidentally deleted. A System Administrator is restoring the table from the most recent snapshot. The snapshot contains all rows that were in the table before the deletion. What is the SIMPLEST solution to restore the table without impacting users?

- A. Restore the snapshot to a new Amazon Redshift cluster, then UNLOAD the table to Amazon S3. In the original cluster, TRUNCATE the table, then load the data from Amazon S3 by using a COPY command.
- B. Use the Restore Table from a Snapshot command and specify a new table name DROP the original table, then RENAME the new table to the original table name.
- C. Restore the snapshot to a new Amazon Redshift cluster. Create a DBLINK between the two clusters in the original cluster, TRUNCATE the destination table, then use an INSERT command to copy the data from the new cluster.
- D. Use the ALTER TABLE REVERT command and specify a time stamp of immediately before the data deletion. Specify the Amazon Resource Name of the snapshot as the SOURCE and use the OVERWRITE REPLACE option.
- E. D. Use the ALTER TABLE REVERT command and specify a time stamp of immediately before the data deletion.

Specify the Amazon Resource Name of the snapshot as the SOURCE and use the OVERWRITE REPLACEMENT option.

Answer: B

Question: 80

An organization's data warehouse contains sales data for reporting purposes. Data governance policies prohibit staff from accessing the customers' credit card numbers. How can these policies be adhered to and still allow a Data Scientist to group transactions that use the same credit card number?

- A. Store a cryptographic hash of the credit card number.
- B. Encrypt the credit card number with a symmetric encryption key, and give the key only to the authorized Data Scientist.
- C. Mask the credit card numbers to only show the last four digits of the credit card number.
- D. Encrypt the credit card number with an asymmetric encryption key and give the decryption key only to the authorized Data Scientist

Answer: C

Question: 81

A real-time bidding company is rebuilding their monolithic application and is focusing on serving real-time data. A large number of reads and writes are generated from thousands of concurrent users who follow items and bid on the company's sale offers. The company is experiencing high latency during special event spikes, with millions of concurrent users. The company needs to analyze and aggregate a part of the data in near real time to feed an internal dashboard. What is the BEST approach for serving and analyzing data, considering the constraint of the row latency on the highly demanded data?

- A. Use Amazon Aurora with Multi Availability Zone and read replicas. Use Amazon ElastiCache in front of the read replicas to serve read-only content quickly. Use the same database as data source for the dashboard.
- B. Use Amazon DynamoDB to store real-time data with Amazon DynamoDB Accelerator to serve content quickly. Use Amazon DynamoDB Streams to replay all changes to the table, process and stream to Amazon ElastiCache with AWS Lambda.
- C. Use Amazon RDS with Multi Availability Zone. Provisioned IOPS EBS volume for storage. Enable up to five read replicas to serve read-only content quickly. Use Amazon EMR with Sqoop to import Amazon RDS data into HDFS for analysis.
- D. Use Amazon Redshift with a DC2 node type and a multi-mode cluster. Create an Amazon EC2 instance with pgpool1 installed. Create an Amazon ElastiCache cluster and route read requests through pgpool1, and use Amazon Redshift for analysis. D

Answer: See Explanation

Question: 82

A gas company needs to monitor gas pressure in their pipelines. Pressure data is streamed from sensors placed throughout the pipelines to monitor the data in real time. When an anomaly is detected, the system must send a notification to open valve. An Amazon Kinesis stream collects the data from the sensors and an anomaly Kinesis

stream triggers an AWS Lambda function to open the appropriate valve. Which solution is the MOST cost-effective for responding to anomalies in real time?

- A. Attach a Kinesis Firehose to the stream and persist the sensor data in an Amazon S3 bucket. Schedule an AWS Lambda function to run a query in Amazon Athena against the data in Amazon S3 to identify anomalies. When a change is detected, the Lambda function sends a message to the anomaly stream to open the valve.
- B. Launch an Amazon EMR cluster that uses Spark Streaming to connect to the Kinesis stream and Spark machine learning to detect anomalies. When a change is detected, the Spark application sends a message to the anomaly stream to open the valve.
- C. Launch a fleet of Amazon EC2 instances with a Kinesis Client Library application that consumes the stream and aggregates sensor data over time to identify anomalies. When an anomaly is detected, the application sends a message to the anomaly stream to open the valve.
- D. Create a Kinesis Analytics application by using the RANDOM_CUT_FOREST function to detect an anomaly. When the anomaly score that is returned from the function is outside of an acceptable range, a message is sent to the anomaly stream to open the valve.

Answer: A

Question: 83

A gaming organization is developing a new game and would like to offer real-time competition to their users. The data architecture has the following characteristics: ? The game application is writing events directly to Amazon DynamoDB from the user's mobile device. ? Users from the website can access their statistics directly from DynamoDB. ? The game servers are accessing DynamoDB to update the user's information. ? The data science team extracts data from DynamoDB for various applications. The engineering team has already agreed to the IAM roles and policies to use for the data science team and the application. Which actions will provide the MOST security, while maintaining the necessary access to the website and game application? (Choose two.)

- A. Use Amazon Cognito user pool to authenticate to both the website and the game application.
- B. Use IAM identity federation to authenticate to both the website and the game application.
- C. Create an IAM policy with PUT permission for both the website and the game application.
- D. Create an IAM policy with fine-grained permission for both the website and the game application.
- E. Create an IAM policy with PUT permission for the game application and an IAM policy with GET permission for the website.

Answer: B,E

Question: 84

An organization has 10,000 devices that generate 100 GB of telemetry data per day, with each record size around 10 KB. Each record has 100 fields, and one field consists of unstructured log data with a "String" data type in the English language. Some fields are required for the real-time dashboard, but all fields must be available for long-term generation. The organization also has 10 PB of previously cleaned and structured data, partitioned by Date, in a SAN that must be migrated to AWS within one month. Currently, the organization does not have any real-time capabilities in their solution. Because of storage limitations in the on-premises data warehouse, selective data is loaded while generating the long-term trend with ANSI SQL queries through JDBC for visualization. In addition to the one-time data loading, the organization needs a cost-effective and real-time solution. How can these requirements be met? (Choose two.)

- A. use AWS IoT to send data from devices to an Amazon SQS queue, create a set of workers in an Auto Scaling group and read records in batch from the queue to process and save the data. Fan out to an Amazon SNS queue attached with an AWS Lambda function to filter the request dataset and save it to Amazon Elasticsearch Service for real-time analytics.
- B. D. Use AWS IoT to send the data from devices to Amazon Kinesis Data Streams with the IoT rules engine. Use one Kinesis Data Firehose stream attached to a Kinesis stream to stream the data into an Amazon S3 bucket partitioned by date. Attach an AWS Lambda function with the same Kinesis stream to filter out the required fields for ingestion into Amazon DynamoDB for real-time analytics.
- C. Use AWS IoT to send the data from devices to Amazon Kinesis Data Streams with the IoT rules engine. Use one Kinesis Data Firehose stream attached to a Kinesis stream to batch and stream the data partitioned by date. Use another Kinesis Firehose stream attached to the same Kinesis stream to filter out the required fields to ingest into Elasticsearch for real-time analytics.
- D. Use AWS IoT to send the data from devices to Amazon Kinesis Data Streams with the IoT rules engine. Use one Kinesis Data Firehose stream attached to a Kinesis stream to stream the data into an Amazon S3 bucket partitioned by date. Attach an AWS Lambda function with the same Kinesis stream to filter out the required fields for ingestion into Amazon DynamoDB for real-time analytics.
- E. use multiple AWS Snowball Edge devices to transfer data to Amazon S3, and use Amazon Athena to query the data.

Answer: A,D

Question: 85

An organization is designing a public web application and has a requirement that states all application users must be centrally authenticated before any operations are permitted. The organization will need to create a user table with fast data lookup for the application in which a user can read only his or her own data. All users already have an account with amazon.com. How can these requirements be met?

- A. Create an Amazon RDS Aurora table, with Amazon_ID as the primary key. The application uses amazon.com web identity federation to get a token that is used to assume an IAM role from AWS STS. Use IAM database authentication by using the rds:db-tag IAM authentication policy and GRANT Amazon RDS row-level read permission per user.
- B. D. Create an Amazon DynamoDtable, with Amazon_ID as the partition key. The application uses amazon.com web identity federation to assume an IAM role from AWS STS in the Role, use IAM condition context key dynamodb:LeadingKeys with IAM substitution variables \$ {www.amazon.com:user_id} and allow the required DynamoDAPI operations in IAM JSON policy Action element for reading the records.
- C. Create an Amazon DynamoDB table, with Amazon_ID as the partition key. The application uses amazon.com web identity federation to get a token that is used to assume an IAM role from AWS STS in the Role, use IAM condition context key dynamodb:LeadingKeys with IAM substitution variables \$ and allow the required DynamoDB API operations in IAM JSON policy Action element for reading the records.
{www.amazon.com:user_id}
- D. D. Create an Amazon DynamoDB table, with Amazon_Ias the partition key. The application uses amazon.com web identity federation to assume an IAM role from AWS STS in the Role, use IAM condition context key dynamodb:LeadingKeys with IAM substitution variables \$ {www.amazon.com:user_id} and allow the required DynamoDB API operations in IAM JSON policy Action element for reading the records.

Answer: C

For More exams visit <https://killexams.com/vendors-exam-list>



Kill your exam at First Attempt....Guaranteed!