DUMPSARENA

AWS Certified Big Data - Specialty (BDS-C00)

Amazon AWS AWS-Certified-Big-Data-Specialty-BDS-C00

Version Demo

Total Demo Questions: 15

Total Premium Questions: 264

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Which of the following are true regarding AWS Cloud Trail?

Choose 3 answers

A. Cloudtrail is enabled globally

B. Cloudtrail is enabled by default

C. Cloudtrail is enabled on a per-region basis

D. Cloudtrail is enabled on a per-service basis

E. Logs can be delivered to a single Amazon S3 bucket for aggregation

F. Logs can only be processes and delivered to the region in which they are generated

ANSWER: A C E

QUESTION NO: 2

What does the "Server Side Encryption" option on Amazon S3 provide?

A. It provides an encrypted virtual disk in the Cloud.

B. It doesn't exist for Amazon S3, but only for Amazon EC2.

C. It encrypts the files that you send to Amazon S3, on the server side.

D. It allows to upload files using an SSL endpoint, for a secure transfer.

ANSWER: A

QUESTION NO: 3

HTTP Query-based requests are HTTP requests that use the HTTP verb GET or POST and a Query parameter named .

A. Action

B. Value

C. Reset

D. Retrieve

ANSWER: A

QUESTION NO: 4

When an Auto Scaling group is running in Amazon Elastic Compute Cloud (EC2), your application rapidly scales up and down in response to load within a 10-minutes window; however, after the load peaks, you begin to see problems in your configuration management system where previously terminated Amazon EC2 resources are still showing as active.

What would be a reliable and efficient way to handle the cleanup of Amazon EC2 resources with your configuration management systems?

Choose 2 answers

A. Write a script that is run by a daily cron job on an Amazon EC2 instance and that executes API Describe calls of the EC2 Auto Scaling group and removes terminated instances from the configuration management system

B. Configure an Amazon Simple Queue Service (SQS) queue for Auto Scaling actions that has a script that listens for new messages and removes terminated instances from the configuration management system

C. Use your existing configuration management system to control the launching and bootstrapping of instances to reduce the number of moving parts in the automation

D. Write a small script that is run during Amazon EC2 instance shutdown to de-register the resource from

the configuration management system

E. Use Amazon Simple Workflow Service (SWF) to maintain an Amazon DynamoDB database that contains a whitelist of instances that have been previously launched, and allow the Amazon SWF worker to remove information from the configuration management system

ANSWER: A D

QUESTION NO: 5

You have an application running on an Amazon Elastic Compute Cloud instance, that uploads 5 GB video objects to Amazon Simple Storage Service (S3). Video uploads are taking longer than expected, resulting in poor application performance. Which method will help improve performance of your application?

A. Enable enhanced networking

B. Use Amazon S3 multipart upload

C. Leveraging Amazon CloudFront, use the HTTP POST method to reduce latency.

D. Use Amazon Elastic Block Store Provisioned IOPs and use an Amazon EBS-optimized instance

ANSWER: B

QUESTION NO: 6

An online photo album app has a key design feature to support multiple screens (e.g, desktop, mobile phone, and tablet) with high-quality displays. Multiple versions of the image must be saved in different resolutions and layouts.

The image-processing Java program takes an average of five seconds per upload, depending on the image size and format. Each image upload captures the following image metadata: user, album, photo label, upload timestamp.

The app should support the following requirements:

Hundreds of user image uploads per second Maximum image upload size of 10 MB Maximum image metadata size of 1 KB

Image displayed in optimized resolution in all supported screens no later than one minute after image upload Which strategy should be used to meet these requirements?

A. Write images and metadata to Amazon Kinesis. Use a Kinesis Client Library (KCL) application to run the image processing and save the image output to Amazon S3 and metadata to the app repository DB.

B. Write image and metadata RDS with BLOB data type. Use AWS Data Pipeline to run the image processing and save the image output to Amazon S3 and metadata to the app repository DB.

C. Upload image with metadata to Amazon S3, use Lambda function to run the image processing and save the images output to Amazon S3 and metadata to the app repository DB.

D. Write image and metadata to Amazon Kinesis. Use Amazon Elastic MapReduce (EMR) with Spark Streaming to run image processing and save the images output to Amazon S3 and metadata to app repository DB.

ANSWER: C

Explanation:

https://aws.amazon.com/blogs/big-data/building-and-maintaining-an-amazon-s3-metadata-index-without-servers

QUESTION NO: 7

A company is using Amazon Machine Learning as part of a medical software application. The application will predict the most likely blood type for a patient based on a variety of other clinical tests that are available when blood type knowledge is unavailable.

What is the appropriate model choice and target attribute combination for the problem?

A. Multi-class classification model with a categorical target attribute

B. Regression model with a numeric target attribute

C. Binary Classification with a categorical target attribute

D. K-Nearest Neighbors model with a multi-class target attribute

ANSWER: A

QUESTION NO: 8

Which of the following notification endpoints or clients are supported by Amazon Simple

Notification Service? Choose 2 answers

A. Email

B. CloudFront distribution

C. File Transfer Protocol

D. Short Message Service

E. Simple Network Management Protocol

ANSWER: B C

QUESTION NO: 9

Customers have recently been complaining that your web application has randomly stopped responding.

During a deep dive of your logs, the team has discovered a major bug in your Java web application. This bug

is causing a memory leak that eventually causes the application to crash.

Your web application runs on Amazon EC2 and was built with AWS CloudFormation.

Which techniques should you see to help detect theses problems faster, as well as help eliminate the server's unresponsiveness?

Choose 2 answers

A. Update your AWS CloudFormation configuration and enable a CustomResource that uses cfn- signal to detect memory leaks

B. Update your CloudWatch metric granularity config for all Amazon EC2 memory metrics to support five-second granularity. Create a CloudWatch alarm that triggers an Amazon SNS notification to page your team when the application memory becomes too large

C. Update your AWS CloudFormation configuration to take advantage of Auto Scaling groups.

Configure an Auto Scaling group policy to trigger off your custom CloudWatch metrics

D. Create a custom CloudWatch metric that you push your JVM memory usage to create a CloudWatch alarm that triggers an Amazon SNS notification to page your team when the application memory usage becomes too large

E. Update your AWS CloudFormation configuration to take advantage of CloudWatch metrics Agent. Configure the CloudWatch Metrics Agent to monitor memory usage and trigger an Amazon SNS alarm

ANSWER: C D

QUESTION NO: 10

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. Use the EMR File System and Amazon S3 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: B

QUESTION NO: 11

Which of the following instance types are available as Amazon EBS backend only?

A. General purpose T2

B. General purpose M3

C. Compute-optimized C4

D. Compute-optimized C3

E. Storage-optimized 12

ANSWER: A C

QUESTION NO: 12

You have been asked to handle a large data migration from multiple Amazon RDS MySQL instances to a DynamoDB table. You have been given a short amount of time to complete the data migration. What will allow you to complete this complex data processing workflow?

A. Create an Amazon Kinesis data stream, pipe in all of the Amazon RDS data, and direct data toward DynamoDB table

B. Write a script in you language of choice, install the script on an Amazon EC2 instance, and then use Auto Scaling groups to ensure that the latency of the mitigation pipelines never exceeds four seconds in any 15-minutes period.

C. Write a bash script to run on your Amazon RDS instance that will export data into DynamoDB

D. Create a data pipeline to export Amazon RDS data and import the data into DynamoDB

ANSWER: D

QUESTION NO: 13

An administrator needs to design a distribution strategy for a star schema in a Redshift cluster. The administrator needs to determine the optimal distribution style for the tables in the Redshift schema.

In which three circumstances would choosing Key-based distribution be most appropriate? (Select three)

A. When the administrator needs to optimize a large, slowly changing dimension table

B. When the administrator needs to reduce cross-node traffic

C. When the administrator needs to optimize the fact table for parity with the number of slices

D. When the administrator needs to balance data distribution and collocation of data

E. When the administrator needs to take advantage of data locality on a local node of joins and aggregates

ANSWER: B C E

QUESTION NO: 14

An organization uses a custom map reduce application to build monthly reports based on many small data files in an Amazon S3 bucket. The data is submitted from various business units on a frequent but unpredictable

schedule. As the dataset continues to grow, it becomes increasingly difficult to process all of the data in one day. The organization has scaled up its Amazon EMR cluster, but other optimizations could improve performance.

The organization needs to improve performance minimal changes to existing processes and applications. What action should the organization take?

A. Use Amazon S3 Event Notifications and AWS Lambda to create a quick search file index in DynamoDB.

B. Add Spark to the Amazon EMR cluster and utilize Resilient Distributed Datasets in-memory.

C. Use Amazon S3 Event Notifications and AWS Lambda to index each file into an Amazon Elasticsearch Service cluster.

D. Schedule a daily AWS Data Pipeline process that aggregates content into larger files using S3DistCp.

E. Have business units submit data via Amazon Kinesis Firehose to aggregate data hourly into Amazon S3.

ANSWER: D

Explanation:

https://aws.amazon.com/blogs/big-data/seven-tips-for-using-s3distcp-on-amazon-emr-to-move-data-efficiently-b

QUESTION NO: 15

Your company operates a website for promoters to sell tickets for entertainment events. You are using a load balancer in front of an Auto Scaling group of web server. Promotion of popular events can cause surges of websites visitors. During scaling-out at theses times, newly launched instances are unable to complete configuration quickly enough, leading to user disappointment.

What option should you choose to improve scaling yet minimize costs? Choose 2 answers

A. Create an AMI with the application pre-configured. Create a new Auto Scaling launch configuration using this new AMI, and configure the Auto Scaling group to launch with this AMI

B. Use Auto Scaling pre-warming to launch instances before they are required. Configure pre- warming to use the CPU trend CloudWatch metric for the group

C. Publish a custom CloudWatch metric from your application on the number of tickets sold, and create an Auto Scaling policy based on this

D. Using the history of past scaling events for similar event sales to predict future scaling requirements. Use the Auto Scaling scheduled scaling feature to vary the size of the fleet

E. Configure an Amazon S3 bucket for website hosting. Upload into the bucket an HTML holding page with its 'x-amz-website-redirect-location' metadata property set to the load balancer endpoint. Configure Elastic Load Balancing to redirect to the holding page when the load on web servers is above a certain level

ANSWER: D E