

DUMPSBOSS.

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://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co

dumpsboss.co

QUESTION NO: 1

You have been asked to use your department's existing continuous integration (CI) tool to test a three- tier

web architecture defined in an AWS CloudFormation template. The tool already supports AWS APIs and can launch new AWS CloudFormation stacks after polling version control. The CI tool reports on the success of the AWS CloudFormation stack creation by using the DescribeStacks API to look for the CREATE_COMPLETE status.

The architecture tiers defined in the template consist of:

- . One load balancer
 - . Five Amazon EC2 instances running the web application
 - . One multi-AZ Amazon RDS instance
- How would you implement this? Choose 2 answers

A. Define a WaitCondition and a WaitConditionHandle for the output of a output of a UserData command that does sanity checking of the application's post-install state

B. Define a CustomResource and write a script that runs architecture-level integration tests through the

load balancer to the application and database for the state of multiple tiers

C. Define a WaitCondition and use a WaitConditionHandle that leverages the AWS SDK to run the DescribeStacks API call until the CREATE_COMPLETE status is returned

D. Define a CustomResource that leverages the AWS SDK to run the DescribeStacks API call until the CREATE_COMPLETE status is returned

E. Define a UserDataHandle for the output of a UserData command that does sanity checking of the application's post-install state and runs integration tests on the state of multiple tiers through load balancer to the application

F. Define a UserDataHandle for the output of a CustomResource that does sanity checking of the application's post-install state

ANSWER: A F

QUESTION NO: 2

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: 3

Does Amazon RDS allow direct host access via Telnet, Secure Shell (SSH), or Windows Remote Desktop Connection?

- A. Yes
- B. No
- C. Depends on if it is in VPC or not

ANSWER: B

QUESTION NO: 4

When attached to an Amazon VPC which two components provide connectivity with external networks?

Choose 2 answers

- A. Elastic IPS (EIP)
- B. NAT Gateway (NAT)
- C. Internet Gateway (IGW)
- D. Virtual Private Gateway (VGW)

ANSWER: C D

QUESTION NO: 5

You have a large number of web servers in an Auto Scaling group behind a load balancer. On an hourly basis, you want to filter and process the logs to collect data on unique visitors, and then put that data in a durable data store in order to run reports. Web servers in the Auto Scaling group are constantly launching and terminating based on your scaling policies, but you do not want to lose any of the log data from these servers during a stop/termination initiated by a user or by Auto Scaling.

What two approaches will meet these requirements? Choose 2 answers

A. Install an Amazon CloudWatch Logs Agent on every web server during the bootstrap process.

Create a CloudWatch log group and define metric Filters to create custom metrics that track unique visitors from the streaming web server logs. Create a scheduled task on an Amazon EC2 instance that runs every hour to generate a new report based on the CloudWatch custom metrics

B. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to Amazon Glacier. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminated. Use Amazon Data pipeline to process data in Amazon Glacier and run reports every hour

C. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to an Amazon S3 bucket. Ensure that the operating system shutdown process triggers a logs transmission when the Amazon EC2 instance is stopped/terminated. Use AWS Data Pipeline to move log data from the Amazon S3 bucket to Amazon Redshift in order to process and run reports every hour

D. Install an AWS Data Pipeline Logs Agent on every web server during the bootstrap process. Create a log group object in AWS Data Pipeline, and define Metric filters to move processed log data directly from the web servers to Amazon Redshift and runs reports every hour

ANSWER: A C

QUESTION NO: 6

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: 7

Is there a method in the IAM system to allow or deny access to a specific instance?

- A.** Only for VPC based instances
- B.** Yes
- C.** No

ANSWER: C

QUESTION NO: 8

You need to configure an Amazon S3 bucket to serve static assets for your public-facing web application.

Which methods ensure that all objects uploaded to the bucket are set to public read? Choose 2 answers

- A. Set permissions on the object to public read during upload
- B. Configure the bucket ACL to sell all objects to public read
- C. Configure the bucket policy to set all objects to public read
- D. Use AWS identity and access Management roles to set the bucket to public read
- E. Amazon S3 objects default to public read, so no action is needed

ANSWER: B C

QUESTION NO: 9

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 these 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

QUESTION NO: 10

What does Amazon CloudFormation provide?

- A. None of these.
- B. The ability to setup Autoscaling for Amazon EC2 instances.
- C. A template to map network resources for Amazon Web Services.
- D. A templated resource creation for Amazon Web Services.

ANSWER: D

QUESTION NO: 11

If I want my instance to run on a single-tenant hardware, which value do I have to set the instance's tenancy attribute to?

- A. dedicated
- B. isolated
- C. one
- D. reserved

ANSWER: A

QUESTION NO: 12

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 metadata 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 image 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 image output to Amazon

ANSWER: D

QUESTION NO: 13

Does AWS Direct Connect allow you access to all Availabilities Zones within a Region?

A. Depends on the type of connection

B. No

C. Yes

D. Only when there's just one availability zone in a region. If there are more than one, only one availability zone can be accessed directly.

ANSWER: A

QUESTION NO: 14

Which DNS name can only be resolved within Amazon EC2?

A. Internal DNS name

B. External DNS name

C. Global DNS name

D. Private DNS name

ANSWER: A

QUESTION NO: 15

A company generates a large number of files each month and needs to use AWS import/export to move these files into Amazon S3 storage. To satisfy the auditors, the company needs to keep a record of which files were imported into Amazon S3.

What is a low-cost way to create a unique log for each import job?

A. Use the same log file prefix in the import/export manifest files to create a versioned log file in Amazon S3 for all imports

B. Use the log file prefix in the import/export manifest file to create a unique log file in Amazon S3 for each import

- C. Use the log file checksum in the import/export manifest file to create a log file in Amazon S3 for each import
- D. Use script to iterate over files in Amazon S3 to generate a log after each import/export job

ANSWER: B

Explanation:

<https://aws.amazon.com/cn/blogs/aws/send-us-that-data/>