

# Amazon

## BDS-C00 Exam

**Amazon AWS Certified Big Data - Speciality Exam**

**Questions & Answers  
Demo**

## Version: 9.0

---

### Question: 1

---

You are currently hosting multiple applications in a VPC and have logged numerous port scans coming in from a specific IP address block. Your security team has requested that all access from the offending IP address block be denied for the next 24 hours.

Which of the following is the best method to quickly and temporarily deny access from the specified IP address block?

- A. Create an AD policy to modify Windows Firewall settings on all hosts in the VPC to deny access from the IP address block
- B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP address block
- C. Add a rule to all of the VPC 5 Security Groups to deny access from the IP address block
- D. Modify the Windows Firewall settings on all Amazon Machine Images (AMIs) that your organization uses in that VPC to deny access from the IP address block

---

**Answer: B**

---

---

### Question: 2

---

The operations team and the development team want a single place to view both operating system and application logs.

How should you implement this using AWS services? Choose two answers

- A. Using AWS CloudFormation, create a CloudWatch Logs LogGroup and send the operating system and application logs of interest using the CloudWatch Logs Agent
- B. Using AWS CloudFormation and configuration management, set up remote logging to send events via UDP packets to CloudTrail
- C. Using configuration management, set up remote logging to send events to Amazon Kinesis and insert these into Amazon CloudSearch or Amazon Redshift, depending on available analytic tools
- D. Using AWS CloudFormation, create a CloudWatch Logs LogGroup. Because the CloudWatch log agent automatically sends all operating system logs, you only have to configure the application logs for sending off-machine
- E. Using AWS CloudFormation, merge the application logs with the operating system logs, and use IAM Roles to allow both teams to have access to view console output from Amazon EC2

---

**Answer: A, C**

---

---

**Question: 3**

---

You are working with customer who has 10 TB of archival data that they want to migrate to Amazon Glacier. The customer has a 1Mbps connection to the Internet. Which service or feature provide the fastest method of getting the data into Amazon Glacier?

- A. Amazon Glacier multipart upload
- B. AWS Storage Gateway
- C. VM Import/Export
- D. AWS Import/Export

---

**Answer: D**

---

---

**Question: 4**

---

A user has provisioned 2000 IOPS to the EBS volume. The application hosted on that EBS is experiencing less IOPS than provisioned. Which of the below mentioned options does not affect the IOPS of the volume?

- A. The application does not have enough IO for the volume
- B. The instance is EBS optimized
- C. The EC2 instance has 10 Gigabit Network connectivity
- D. The volume size is too large

---

**Answer: D**

---

---

**Question: 5**

---

You want to securely distribute credentials for your Amazon RDS instance to your fleet of web server instances. The credentials are stored in a file that is controlled by a configuration management system. How do you securely deploy the credentials in an automated manner across the fleet of web server instances, which can number in the hundreds, while retaining the ability to roll back if needed?

- A. Store your credential files in an Amazon S3 bucket. Use Amazon S3 server-side encryption on the credential files. Have a scheduled job that pulls down the credential files into the instances every 10 minutes
- B. Store the credential files in your version-controlled repository with the rest of your code. Have a post-commit action in version control that kicks off a job in your continuous integration system which securely copies the new credentials files to all web server instances
- C. Insert credential files into user data and use an instance lifecycle policy to periodically refresh the files from the user data

- D. Keep credential files as a binary blob in an Amazon RDS MySQL DB instance, and have a script on each Amazon EC2 instance that pulls the files down from the RDS instance
- E. Store the credential files in your version-controlled repository with the rest of your code. Use a parallel file copy program to send the credential files from your local machine to the Amazon EC2 instances

---

**Answer: D**

---

---

**Question: 6**

---

A us-based company is expanding their web presence into Europe. The company wants to extend their AWS infrastructure from Northern Virginia (us-east-1) into the Dublin (eu-west-1) region. Which of the following options would enable an equivalent experience for users on both continents?

- A. Use a public-facing load balancer per region to load-balancer web traffic, and enable HTTP health checks
- B. Use a public-facing load balancer per region to load balancer web traffic, and enable sticky sessions
- C. Use Amazon Route S3, and apply a geolocation routing policy to distribution traffic across both regions
- D. Use Amazon Route S3, and apply a weighted routing policy to distribute traffic across both regions

---

**Answer: C**

---

---

**Question: 7**

---

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: 8**

---

You have started a new job and are reviewing your company's infrastructure on AWS You notice one web application where they have an Elastic Load Balancer (&B) in front of web instances in an Auto Scaling Group When you check the metrics for the ELB in CloudWatch you see four healthy instances in Availability Zone (AZ) A and zero in AZ B There are zero unhealthy instances. What do you need to fix to balance the instances across AZs?

- A. Set the ELB to only be attached to another AZ
- B. Make sure Auto Scaling is configured to launch in both AZs
- C. Make sure your AMI is available in both AZs
- D. Make sure the maximum size of the Auto Scaling Group is greater than 4

---

**Answer: B**

---

---

**Question: 9**

---

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: 10**

---

In AWS, which security aspects are the customer's responsibility? Choose 4 answers

- A. Life-Cycle management of IAM credentials
- B. Security Group and ACL settings
- C. Controlling physical access to compute resources
- D. Patch management on the EC2 instance's operating system
- E. Encryption of EBS volumes

F. Decommissioning storage devices

---

**Answer: A, B, D, E**

---

**Thank You For Trying BDS-C00 PDF Demo**

**To try our BDS-C00 Premium Files visit link below:**

**<https://examsland.com/latest-exam-questions/BDS-C00/>**

**Start Your BDS-C00 Preparation**

**Use Coupon **EL25** for extra 25% discount on the purchase of Practice Test Software.**