

DUMPS ARENA

AWS Certified SysOps Administrator - Associate

Amazon AWS AWS-SysOps-SOA-C01

Version Demo

Total Demo Questions: 20

Total Premium Questions: 932

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

A company has deployed a fleet of Amazon EC2 web servers for the upcoming release of a new product. The SysOps Administrator needs to test the Amazon CloudWatch notification settings for this deployment to ensure that a notification is sent using Amazon SNS if the CPU utilization of an EC2 instance exceeds 70%.

How should the Administrator accomplish this?

- A. Use the set-alarm-state command in AWS CloudTrail to invoke the Amazon SNS notification
- B. Use CloudWatch custom metrics to set the alarm state in AWS CloudTrail and enable Amazon SNS notifications
- C. Use EC2 instance metadata to manually set the CPU utilization to 75% and invoke the alarm state
- D. Use the set-alarm-state command in the AWS CLI for CloudWatch

ANSWER: D**QUESTION NO: 2**

A user has launched an EC2 instance and deployed a production application in it. The user wants to prohibit any mistakes from the production team to avoid accidental termination. How can the user achieve this?

- A. The user can the set DisableApiTermination attribute to avoid accidental termination
- B. It is not possible to avoid accidental termination
- C. The user can set the Deletion termination flag to avoid accidental termination
- D. The user can set the InstanceInitiatedShutdownBehavior flag to avoid accidental termination

ANSWER: A**Explanation:**

It is always possible that someone can terminate an EC2 instance using the Amazon EC2 console, command line interface or API by mistake. If the admin wants to prevent the instance from being accidentally terminated, he can enable termination protection for that instance. The DisableApiTermination attribute controls whether the instance can be terminated using the console, CLI or API. By default, termination protection is disabled for an EC2 instance. When it is set it will not allow the user to terminate the instance from CLI, API or the console.

QUESTION NO: 3

A company must ensure that any objects uploaded to an S3 bucket are encrypted.

Which of the following actions will meet this requirement? (Choose two.)

- A. Implement AWS Shield to protect against unencrypted objects stored in S3 buckets.
- B. Implement Object access control list (ACL) to deny unencrypted objects from being uploaded to the S3 bucket.
- C. Implement Amazon S3 default encryption to make sure that any object being uploaded is encrypted before it is stored.
- D. Implement Amazon Inspector to inspect objects uploaded to the S3 bucket to make sure that they are encrypted.
- E. Implement S3 bucket policies to deny unencrypted objects from being uploaded to the buckets.

ANSWER: C E

Explanation:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/dev/bucket-encryption.html>

<https://aws.amazon.com/blogs/security/how-to-prevent-uploads-of-unencrypted-objects-to-amazon-s3/>

QUESTION NO: 4

You have two Elastic Compute Cloud (EC2) instances inside a Virtual Private Cloud (VPC) in the same Availability Zone (AZ) but in different subnets. One instance is running a database and the other instance an application that will interface with the database. You want to confirm that they can talk to each other for your application to work properly. Which two things do we need to confirm in the VPC settings so that these EC2 instances can communicate inside the VPC? (Choose two.)

- A. A network ACL that allows communication between the two subnets.
- B. Both instances are the same instance class and using the same Key-pair.
- C. That the default route is set to a NAT instance or internet Gateway (IGW) for them to communicate.
- D. Security groups are set to allow the application host to talk to the database on the right port/protocol.

ANSWER: A D

QUESTION NO: 5

A SysOps Administrator noticed that the cache hit ratio for an Amazon CloudFront distribution is less than 10%.

Which collection of configuration changes will increase the cache hit ratio for the distribution? (Choose two.)

- A. Ensure that only required cookies, query strings, and headers are forwarded in the Cache Behavior Settings
- B. Change the Viewer Protocol Policy to use HTTPS only
- C. Configure the distribution to use presigned cookies and URLs to restrict access to the distribution
- D. Enable automatic compression of objects in the Cache Behavior Settings

E. Increase the CloudFront time to live (TTL) settings in the Cache Behavior Settings

ANSWER: A E

QUESTION NO: 6

A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console?

- A. The IP address may be attached to one of the instances
- B. The IP address belongs to a different zone than the subnet zone
- C. The user has not created an internet gateway
- D. The IP addresses belong to EC2 Classic; so they cannot be assigned to VPC

ANSWER: D

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP, then it will only have a private IP when launched. If the user wants to connect to an instance from the Internet, he should create an elastic IP with VPC. If the elastic IP is a part of EC2 Classic, it cannot be assigned to a VPC instance.

QUESTION NO: 7

A company needs to deploy a web application on two Amazon EC2 instances behind an Application Load Balancer (ALB). Two EC2 instances will also be deployed to host the database. The infrastructure needs to be designed across Availability Zones (AZs) for high availability and must limit public access to the instances as much as possible.

How should this be achieved within a VPC?

- A. Use two AZs and create a public subnet in each AZ for the Application Load Balancer, a private subnet in each AZ for the web servers, and a private subnet in each AZ for the database servers.
- B. Use two AZs and create a public subnet in each AZ for the Application Load Balancer, a public subnet in each AZ for the web servers, and a public subnet in each AZ for the database servers.
- C. Use two AZs and create one public subnet for the Application Load Balancer, a private subnet in each AZ for the web servers, and a public subnet in each AZ for the database servers.
- D. Use two AZs and create one public subnet for the Application Load Balancer, a public subnet in each AZ for the web servers, and a private subnet in each AZ for the database servers.

ANSWER: A

QUESTION NO: 8

Your company is moving towards tracking web page users with a small tracking

image loaded on each page. Currently, you are serving this image out of US-East, but you are starting to get concerned about the time it takes to load the image for users on the west coast.

What are the two best ways to speed up serving this image? (Choose two.)

- A. Use Route 53's Latency Based Routing and serve the image out of US-West-2 as well as US-East-1
- B. Serve the image out through CloudFront
- C. Serve the image out of S3 so that it isn't being served out of your web application tier
- D. Use EBS PIOPs to serve the image faster out of your EC2 instances

ANSWER: A B

Explanation:

CloudFront gets the image closer to the user and Route53 ensures the best connection based on network latency.

QUESTION NO: 9

Your organization's security policy requires that all privileged users either use frequently rotated passwords or one-time access credentials in addition to username/password.

Which two of the following options would allow an organization to enforce this policy for AWS users? (Choose two.)

- A. Configure multi-factor authentication for privileged IAM users
- B. Create IAM users for privileged accounts
- C. Implement identity federation between your organization's Identity provider leveraging the IAM Security Token Service
- D. Enable the IAM single-use password policy option for privileged users

ANSWER: A B

QUESTION NO: 10

What does the "configure" command allow an Administrator to do when setting up the AWS CLI? (Choose two.)

- A. Decide which VPC to create instances in.

- B. Designate the format of the response to CLI commands.
- C. Choose the default EC2 instance.
- D. Encrypt the CLI commands.
- E. Designate the default region.

ANSWER: B E

QUESTION NO: 11

A sysops administrator has an AWS Lambda function that performs maintenance on various AWS resources. This function must be run nightly.

Which is the MOST cost-effective solution?

- A. Launch a single t2.nano Amazon EC2 instance and create a Linux cron job to invoke the Lambda function at the same time every night.
- B. Set up an Amazon CloudWatch metrics alarm to invoke the Lambda function at the same time every night.
- C. Schedule an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function at the same time every night.
- D. Implement a Chef recipe in AWS OpsWorks stack to invoke the Lambda function at the same time every night.

ANSWER: C

Explanation:

Reference: <https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/RunLambdaSchedule.html>

QUESTION NO: 12

A user has created a VPC with the public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The public subnet uses CIDR 20.0.1.0/24. The user is planning to host a web server in the public subnet (port 80) and a DB server in the private subnet (port 3306). The user is configuring a security group for the public subnet (WebSecGrp) and the private subnet (DBSecGrp). Which of the below mentioned entries is required in the web server security group (WebSecGrp)?

- A. Configure Destination as DB Security group ID (DbSecGrp) for port 3306 Outbound
- B. 80 for Destination 0.0.0.0/0 Outbound
- C. Configure port 3306 for source 20.0.0.0/24 InBound
- D. Configure port 80 InBound for source 20.0.0.0/16

ANSWER: A**Explanation:**

A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet to host the web server and DB server respectively, the user should configure that the instances in the public subnet can receive inbound traffic directly from the internet. Thus, the user should configure port 80 with source 0.0.0.0/0 in InBound. The user should configure that the instance in the public subnet can send traffic to the private subnet instances on the DB port. Thus, the user should configure the DB security group of the private subnet (DbSecGrp) as the destination for port 3306 in Outbound.

QUESTION NO: 13

Which of the following requires a custom CloudWatch metric to monitor?

- A. Data transfer of an EC2 instance
- B. Disk usage activity of an EC2 instance
- C. Memory Utilization of an EC2 instance
- D. CPU Utilization of an EC2 instance

ANSWER: C**Explanation:**

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/ec2-metricscollected.html> CPU, Disk I/O, Data Transfer are default metrics. Memory is not mentioned.

QUESTION NO: 14

An application accesses data through a file system interface. The application runs on Amazon EC2 instances in multiple Availability Zones, all of which must share the same data. While the amount of data is currently small, the company anticipates that it will grow to tens of terabytes over the lifetime of the application.

What is the MOST scalable storage solution to fulfill the requirement?

- A. Connect a large Amazon EBS volume to multiple instances and schedule snapshots.
- B. Deploy Amazon EFS in the VPC and create mount targets in multiple subnets.
- C. Launch an EC2 instance and share data using SMB/CIFS or NFS.
- D. Deploy an AWS Storage Gateway cached volume on Amazon EC2.

ANSWER: D

QUESTION NO: 15

A user is accessing RDS from an application. The user has enabled the Multi AZ feature with the MS SQL RDS DB. During a planned outage how will AWS ensure that a switch from DB to a standby replica will not affect access to the application?

- A. RDS will have an internal IP which will redirect all requests to the new DB
- B. RDS uses DNS to switch over to stand by replica for seamless transition
- C. The switch over changes Hardware so RDS does not need to worry about access
- D. RDS will have both the DBs running independently and the user has to manually switch over

ANSWER: B**Explanation:**

In the event of a planned or unplanned outage of a DB instance, Amazon RDS automatically switches to a standby replica in another Availability Zone if the user has enabled Multi AZ. The automatic failover mechanism simply changes the DNS record of the DB instance to point to the standby DB instance. As a result, the user will need to re-establish any existing connections to the DB instance. However, as the DNS is the same, the application can access DB seamlessly.

QUESTION NO: 16

According to the shared responsibility model, for which of the following Amazon EC2 activities is AWS responsible? (Choose two.)

- A. Patching the guest operating system
- B. Monitoring memory utilization
- C. Configuring network ACLs
- D. Patching the hypervisor
- E. Maintaining network infrastructure

ANSWER: D E**Explanation:**

Reference: <https://aws.amazon.com/compliance/shared-responsibility-model/>

QUESTION NO: 17

What are characteristics of Amazon S3? (Choose two.)

- A. Objects are directly accessible via a URL
- B. S3 should be used to host a relational database

- C. S3 allows you to store objects or virtually unlimited size
- D. S3 allows you to store virtually unlimited amounts of data
- E. S3 offers Provisioned IOPS

ANSWER: A D

Explanation:

The total volume of data and number of objects you can store are unlimited. Individual Amazon S3 objects can range in size from a minimum of 0 bytes to a maximum of 5 terabytes. The largest object that can be uploaded in a single PUT is 5 gigabytes. For objects larger than 100 megabytes, customers should consider using the Multipart Upload capability. Reference: <https://aws.amazon.com/s3/faqs/>

QUESTION NO: 18

A SysOps Administrator is attempting to use AWS Systems Manager Session Manager to initiate a SSH session with an Amazon EC2 instance running on a custom Linux Amazon Machine Image (AMI). The Administrator cannot find the target instance in the Session Manager console.

Which combination of actions will solve this issue? (Choose two.)

- A. Add Systems Manager permissions to the instance profile.
- B. Configure the bucket used by Session Manager logs to allow write access.
- C. Install Systems Manager Agent on the instance.
- D. Modify the instance security group to allow inbound traffic on SSH port 22.
- E. Reboot the instance with a new SSH key pair named ssm-user.

ANSWER: B D

QUESTION NO: 19

Which of the following are characteristics of Amazon VPC subnets? (Choose two.)

- A. Each subnet maps to a single Availability Zone
- B. A CIDR block mask of /25 is the smallest range supported
- C. Instances in a private subnet can communicate with the internet only if they have an Elastic IP.
- D. By default, all subnets can route between each other, whether they are private or public
- E. V Each subnet spans at least 2 Availability zones to provide a high-availability environment

ANSWER: A D**Explanation:**

“Each subnet must reside entirely within one Availability Zone and cannot span zones.”

“Every subnet that you create is automatically associated with the main route table for the VPC.” Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION NO: 20

Which of the following services is offered by CloudWatch?

- A. Fixing broken links on the client's instances
- B. Creating IAM users for all services in AWS
- C. Monitoring estimated AWS charges
- D. Balancing the request load between various instances

ANSWER: C**Explanation:**

AWS CloudWatch supports monitoring of the AWS estimated usage charges. You create an Amazon CloudWatch alarm that will monitor your estimated Amazon Web Services (AWS) charges. When you enable the monitoring of estimated charges for your AWS account, the estimated charges are calculated and sent several times daily to CloudWatch as metric data.

Reference:

http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/monitor_estimated_charges_with_cloudwatch.html