

AWS_SOA-COO Exam

Volume: 758 Questions

Question: 1

A user has created photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly. Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

- A. AWS Elastic Transcoder
- B. AWS Simple Queue Service
- C. AWS Simple Notification Service
- D. AWS Glacier

Answer: B

Explanation:

Amazon Simple Queue Service (SQS) is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. The user can configure SQS, which will decouple the call between the EC2 application and S3. Thus, the application does not keep waiting for S3 to provide the data.

Reference: <http://aws.amazon.com/sqs/faqs/>

Question: 2

A user has created an EBS volume of 10 GB and attached it to a running instance. The user is trying to access EBS for first time. Which of the below mentioned options is the correct statement with respect to a first time EBS access?

- A. The volume will be blank
- B. If the EBS is mounted it will ask the user to create a file system
- C. The volume will show a loss of the IOPS performance the first time
- D. The volume will show a size of 8 GB

Answer: C

Explanation:

AWS SOA-COO Exam

A user can create an EBS volume either from a snapshot or as a blank volume. If the volume is from a snapshot it will not be blank. The volume shows the right size only as long as it is mounted. This shows that the file system is created. When the user is accessing the volume the AWS EBS will wipe out the block storage or instantiate from the snapshot. Thus, the volume will show a loss of IOPS. It is recommended that the user should pre warm the EBS before use to achieve better IO.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-prewarm.html>

Question: 3

A user has recently started using EC2. The user launched one EC2 instance in the default subnet in EC2-VPC. Which of the below mentioned options is not attached or available with the EC2 instance when it is launched?

- A. Public IP address
- B. Internet gateway
- C. Elastic IP
- D. Private IP address

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to a user's AWS account. A subnet is a range of IP addresses in the VPC. The user can launch the AWS resources into a subnet. There are two supported platforms into which a user can launch instances: EC2-Classic and EC2-VPC (default subnet). A default VPC has all the benefits of EC2-VPC and the ease of use of EC2-Classic. Each instance that the user launches into a default subnet has a private IP address and a public IP address. These instances can communicate with the internet through an internet gateway. An internet gateway enables the EC2 instances to connect to the internet through the Amazon EC2 network edge.

Reference:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.html

Question: 4

A user has created a VPC with CIDR 20.0.0.0/24. The user has created a public subnet with CIDR 20.0.0.0/25. The user is trying to create the private subnet with CIDR 20.0.0.128/25. Which of the below mentioned statements is true in this scenario?

- A. This statement is wrong as AWS does not allow CIDR 20.0.0.0/25

AWS SOA-COO Exam

- B. It will allow the user to create a private subnet with CIDR as 20.0.0.128/25
- C. It will not allow the user to create a private subnet due to a wrong CIDR range
- D. It will not allow the user to create the private subnet due to a CIDR overlap

Answer: B

Explanation:

When the user creates a subnet in VPC, he specifies the CIDR block for the subnet. The CIDR block of a subnet can be the same as the CIDR block for the VPC (for a single subnet in the VPC), or a subset (to enable multiple subnets). If the user creates more than one subnet in a VPC, the CIDR blocks of the subnets must not overlap. Thus, in this case the user has created a VPC with the CIDR block 20.0.0.0/24, which supports 256 IP addresses (20.0.0.0 to 20.0.0.255). The user can break this CIDR block into two subnets, each supporting 128 IP addresses. One subnet uses the CIDR block 20.0.0.0/25 (for addresses 20.0.0.0 - 20.0.0.127) and the other uses the CIDR block 20.0.0.128/25 (for addresses 20.0.0.128 - 20.0.0.255).

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

Question: 5

A user has created a VPC with CIDR 20.0.0.0/24. The user has created a public subnet with CIDR 20.0.0.0/25 and a private subnet with CIDR 20.0.0.128/25. The user has launched one instance each in the private and public subnets. Which of the below mentioned options cannot be the correct IP address (private IP) assigned to an instance in the public or private subnet?

- A. 20.0.0.255
- B. 20.0.0.122
- C. 20.0.0.132
- D. 20.0.0.55

Answer: A

Explanation:

When the user creates a subnet in VPC, he specifies the CIDR block for the subnet. In this case the user has created a VPC with the CIDR block 20.0.0.0/24, which supports 256 IP addresses (20.0.0.0 to 20.0.0.255). The public subnet will have IP addresses between 20.0.0.0 - 20.0.0.127 and the private subnet will have IP addresses between 20.0.0.128 - 20.0.0.255. AWS reserves the first four IP addresses and the last IP address in each subnet's CIDR block. These are not available for the user to use. Thus, the instance cannot have an IP address of 20.0.0.255.

AWS_SOA-COO Exam

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

Question: 6

George has launched three EC2 instances inside the US-East-1a zone with his AWS account. Ray has launched two EC2 instances in the US-East-1a zone with his AWS account. Which of the below mentioned statements will help George and Ray understand the availability zone (AZ) concept better?

- A. The US-East-1a region of George and Ray can be different availability zones
- B. The instances of George and Ray will be running in the same data centre
- C. All the instances of George and Ray can communicate over a private IP without any cost
- D. All the instances of George and Ray can communicate over a private IP with a minimal cost

Answer: A

Explanation:

Each AWS region has multiple, isolated locations known as Availability Zones. To ensure that the AWS resources are distributed across the Availability Zones for a region, AWS independently maps the Availability Zones to identifiers for each account. In this case the Availability Zone US-East-1a where George's EC2 instances are running might not be the same location as the US-East-1a zone of Ray's EC2 instances. There is no way for the user to coordinate the Availability Zones between accounts.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

Question: 7

A user is launching an EC2 instance in the US East region. Which of the below mentioned options is recommended by AWS with respect to the selection of the availability zone?

- A. Always select the US-East-1-a zone for HA
- B. The user can never select the availability zone while launching an instance
- C. Do not select the AZ; instead let AWS select the AZ
- D. Always select the AZ while launching an instance

AWS SOA-COO Exam

Answer: C

Explanation:

When launching an instance with EC2, AWS recommends not to select the availability zone (AZ). AWS specifies that the default Availability Zone should be accepted. This is because it enables AWS to select the best Availability Zone based on the system health and available capacity. If the user launches additional instances, only then an Availability Zone should be specified. This is to specify the same or different AZ from the running instances.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

Question: 8

A user has created a public subnet with VPC and launched an EC2 instance within it. The user is trying to delete the subnet. What will happen in this scenario?

- A. It will delete the subnet and make the EC2 instance as a part of the default subnet
- B. It will not allow the user to delete the subnet until the instances are terminated
- C. The subnet can never be deleted independently, but the user has to delete the VPC first
- D. It will delete the subnet as well as terminate the instances

Answer: B

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 an instance is launched it will have a network interface attached with it. The user cannot delete the subnet until he terminates the instance and deletes the network interface.

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

Question: 9

A user has enabled the Multi AZ feature with the MS SQL RDS database server. Which of the below mentioned statements will help the user understand the Multi AZ feature better?

- A. In a Multi AZ, AWS runs two DBs in parallel and copies the data synchronously to the replica copy
- B. AWS MS SQL does not support the Multi AZ feature

AWS SOA-COO Exam

- C. In a Multi AZ, AWS runs just one DB but copies the data synchronously to the standby replica
- D. In a Multi AZ, AWS runs two DBs in parallel and copies the data asynchronously to the replica copy

Answer: C

Explanation:

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy, eliminate I/O freezes, and minimize latency spikes during system backups. Running a DB instance with high availability can enhance availability during planned system maintenance, and help protect your databases against DB instance failure and Availability Zone disruption.

Note that the high-availability feature is not a scaling solution for read-only scenarios; you cannot use a standby replica to serve read traffic. To service read-only traffic, you should use a read replica.

Reference:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html>

Question: 10

A user has created a queue named “myqueue” with SQS. There are four messages published to queue which are not received by the consumer yet. If the user tries to delete the queue, what will happen?

- A. It will delete the queue
- B. A user can never delete a queue manually. AWS deletes it after 30 days of inactivity on queue
- C. It will ask user to delete the messages first
- D. It will initiate the delete but wait for four days before deleting until all messages are deleted automatically.

Answer: A

Explanation:

SQS allows the user to move data between distributed components of applications so they can perform different tasks without losing messages or requiring each component to be always available. The user can delete a queue at any time, whether it is empty or not. It is important to

AWS SOA-COO Exam

note that queues retain messages for a set period of time. By default, a queue retains messages for four days.

Reference:

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/SQSConcepts.html>

Question: 11

A user is trying to create a PIOPS EBS volume with 3 GB size and 90 IOPS. Will AWS create the volume?

- A. Yes, since the ratio between EBS and IOPS is less than 30
- B. Yes, since PIOPS is higher than 100
- C. No, since the PIOPS and EBS size ratio is less than 30
- D. No, the EBS size is less than 4GB

Answer: D

Explanation:

A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume.

Reference:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_piops

Question: 12

A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created a public subnet CIDR (20.0.0.0/24) and VPN only subnets CIDR (20.0.1.0/24) along with the VPN gateway (vgw-12345) to connect to the user's data centre. The user's data centre has CIDR 172.28.0.0/12. The user has also setup a NAT instance (i-123456) to allow traffic to the internet from the VPN subnet. Which of the below mentioned options is not a valid entry for the main route table in this scenario?

- A. Destination: 172.28.0.0/12 and Target: vgw-12345
- B. Destination: 0.0.0.0/0 and Target: i-12345
- C. Destination: 20.0.0.0/16 and Target: local

AWS SOA-COO Exam

D. Destination: 20.0.1.0/24 and Target: i-12345

Answer: D

Explanation:

The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all traffic of the VPN subnet. If the user has setup a NAT instance to route all the internet requests then all requests to the internet should be routed to it. All requests to the organization's DC will be routed to the VPN gateway. Here are the valid entries for the main route table in this scenario:

Destination: 0.0.0.0/0 & Target: i-12345 (To route all internet traffic to the NAT Instance)

Destination: 172.28.0.0/12 & Target: vgw-12345 (To route all the organization's data centre traffic to the VPN gateway)

Destination: 20.0.0.0/16 & Target: local (To allow local routing in VPC)

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario3.html

Question: 13

A user has a weighing plant. The user measures the weight of some goods every 5 minutes and sends data to AWS CloudWatch for monitoring and tracking. Which of the below mentioned parameters is mandatory for the user to include in the request list?

A. Timezone

B. Value

C. Metric Name

D. Namespace

Answer: D

Explanation:

AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. The user can publish the data to CloudWatch as single data points or as an aggregated set of data points called a statistic set. The user has to always include the namespace as part of the request. The user can supply a file instead of the metric name. If the user does not supply the timezone, it accepts the current time. If the user is sending the data as a single data point it will have parameters, such as value. However, if the user is sending as an aggregate it will have parameters, such as statistic-values.

AWS SOA-COO Exam

Reference:

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/PublishMetrics.html>

Question: 14

A user wants to capture errors that occur in the AWS MySQL RDS DB. Which of the below mentioned activities may help the user to get the data easily?

- A. Direct the error log to a DB table and then query that table
- B. Download the log file to DynamoDB and search for records
- C. Find all the transaction logs and query on those records
- D. It is not possible to get the log files for MySQL RDS

Answer: A

Explanation:

The user can view, download, and watch the database logs using the Amazon RDS console, the Command Line Interface (CLI) or the Amazon RDS API. For the MySQL RDS, the user can view the error log, slow query log, and general logs. The user can also view the MySQL logs easily by directing the logs to a database table in the main database and querying that table.

Reference:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.html

Question: 15

An organization has configured Auto Scaling for hosting their application. The system admin wants to understand the Auto Scaling health check process. If the instance is unhealthy, Auto Scaling launches an instance and terminates the unhealthy instance. What is the order execution?

- A. Auto Scaling launches and terminates the instances simultaneously
- B. Auto Scaling performs the launch and terminate processes in a random order
- C. Auto Scaling terminates the instance first and then launches a new instance
- D. Auto Scaling launches a new instance first and then terminates the unhealthy instance

Answer: C

AWS_SOA-COO Exam

Explanation:

Auto Scaling keeps checking the health of the instances at regular intervals and marks the instance for replacement when it is unhealthy. The ReplaceUnhealthy process terminates instances which are marked as unhealthy and subsequently creates new instances to replace them. This process first terminates the instance and then launches a new instance.

Reference:

http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/US_SuspendResume.html

Question: 16

The CFO of a company wants to allow one of his employees to view only the AWS usage report page. Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

- A. "Effect": "Allow", "Action": ["aws-portal:ViewUsage"], "Resource": "*"
- B. "Effect": "Allow", "Action": ["AccountUsage"], "Resource": "*"
- C. "Effect": "Allow", "Action": ["aws-portal:ViewBilling"], "Resource": "*"
- D. "Effect": "Allow", "Action": ["Describe"], "Resource": "Billing"

Answer: A

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "aws-portal:ViewUsage"
      ],
      "Resource": "*"
    }
  ]
}
```

Reference:

<http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-permissions-ref.html>