

Amazon

Exam Questions AWS-SysOps

Amazon AWS Certified SysOps Administrator - Associate



NEW QUESTION 1

- (Topic 1)

A media company produces new video files on-premises every day with a total size of around 100GBS after compression All files have a size of 1 -2 GB and need to be uploaded to Amazon S3 every night in a fixed time window between 3am and 5am Current upload takes almost 3 hours, although less than half of the available bandwidth is used.

What step(s) would ensure that the file uploads are able to complete in the allotted time window?

- A. Increase your network bandwidth to provide faster throughput to S3
- B. Upload the files in parallel to S3
- C. Pack all files into a single archive, upload it to S3, then extract the files in AWS
- D. Use AWS Import/Export to transfer the video files

Answer: B

Explanation:

Reference:

<http://aws.amazon.com/importexport/faqs/>

NEW QUESTION 2

- (Topic 1)

What would happen to an RDS (Relational Database Service) multi-Availability Zone deployment of the primary OB instance fails?

- A. The IP of the primary DB instance is switched to the standby OB instance
- B. The RDS (Relational Database Service) DB instance reboots
- C. A new DB instance is created in the standby availability zone
- D. The canonical name record (CNAME) is changed from primary to standby

Answer: D

NEW QUESTION 3

- (Topic 1)

Your EC2-Based Multi-tier application includes a monitoring instance that periodically makes application -level read only requests of various application components and if any of those fail more than three times 30 seconds calls CloudWatch to fire an alarm, and the alarm notifies your operations team by email and SMS of a possible application health problem. However, you also need to watch the watcher -the monitoring instance itself - and be notified if it becomes unhealthy.

Which of the following is a simple way to achieve that goal?

- A. Run another monitoring instance that pings the monitoring instance and fires a could watch alarm that notifies your operations team should the primary monitoring instance become unhealth
- B. Set a CloudWatch alarm based on EC2 system and instance status checks and have the alarm notify your operations team of any detected problem with the monitoring instanc
- C. Set a CloudWatch alarm based on the CPU utilization of the monitoring instance and have the alarm notify your operations team if C r the CPU usage exceeds 50% few more than one minute: then have your monitoring application go into a CPU-bound loop should it Detect any application problem
- D. Have the monitoring instances post messages to an SOS queue and then dequeue those messages on another instance should the queue cease to have new messages, the second instance should first terminate the original monitoring instance start anotherbackup monitoring instance and assume (he role of the previous monitoring instance and beginning adding messages to the SQSqueu

Answer: D

NEW QUESTION 4

- (Topic 1)

How can the domain's zone apex for example "myzoneapexdomain com" be pointed towards an Elastic Load Balancer?

- A. By using an AAAA record
- B. By using an A record
- C. By using an Amazon Route 53 CNAME record
- D. By using an Amazon Route 53 Alias record

Answer: D

Explanation:

Reference:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html>

NEW QUESTION 5

- (Topic 1)

You are tasked with setting up a cluster of EC2 Instances for a NoSQL database. The database requires random read IO disk performance up to a 100,000 IOPS at 4KB block size per node.

Which of the following EC2 instances will perform the best for this workload?

- A. A High-Memory Quadruple Extra Large (m2.4xlarge) with EBS-Optimized set to true and a PIOPs EBS volume
- B. A Cluster Compute Eight Extra Large (cc2.8xlarge) using instance storage
- C. High I/O Quadruple Extra Large (hi1.4xlarge) using instance storage
- D. A Cluster GPU Quadruple Extra Large (cg1.4xlarge) using four separate 4000 PIOPS EBS volumes in a RAID 0 configuration

Answer: C

Explanation:

Explanation: Reference:
<http://aws.amazon.com/ec2/instance-types/>

NEW QUESTION 6

- (Topic 1)

A customer has a web application that uses cookie Based sessions to track logged in users It Is deployed on AWS using ELB and Auto Scaling The customer observes that when load increases. Auto Scaling launches new Instances but the load on the easting Instances does not decrease, causing all existing users to have a sluggish experience.

Which two answer choices independently describe a behavior that could be the cause of the sluggish user experience? Choose 2 answers

- A. ELB's normal behavior sends requests from the same user to the same backend instance
- B. ELB's behavior when sticky sessions are enabled causes ELB to send requests in the same session to the same backend instance
- C. A faulty browser is not honoring the TTL of the ELB DNS nam
- D. The web application uses long polling such as comet or websocket
- E. Thereby keeping a connection open to a web server tor a long time
- F. The web application uses long polling such as comet or websocket
- G. Thereby keeping a connection open to a web server for a long tim

Answer: BD

NEW QUESTION 7

- (Topic 1)

You have been asked to leverage Amazon VPC BC2 and SOS to implement an application that submits and receives millions of messages per second to a message queue. You want to ensure your application has sufficient bandwidth between your EC2 instances and SQS Which option will provide the most scalable solution for communicating between the application and SQS?

- A. Ensure the application instances are properly configured with an Elastic Load Balancer
- B. Ensure the application instances are launched in private subnets with the EBS-optimized option enabled
- C. Ensure the application instances are launched in public subnets with the associate-public-IP-address=true option enabled
- D. Launch application instances in private subnets with an Auto Scaling group and Auto Scaling triggers configured to watch the SQS queue size

Answer: B

Explanation:

Reference:
<http://www.cardinalpath.com/autoscaling-your-website-with-amazon-web-services-part-2/>

NEW QUESTION 8

- (Topic 1)

Your entire AWS infrastructure lives inside of one Amazon VPC You have an Infrastructure monitoring application running on an Amazon instance in Availability Zone (AZ) A of the region, and another application instance running in AZ B. The monitoring application needs to make use of ICMP ping to confirm network reachability of the instance hosting the application.

Can you configure the security groups for these instances to only allow the ICMP ping to pass from the monitoring instance to the application instance and nothing else" If so how?

- A. No Two instances in two different AZ's can't talk directly to each other via ICMP ping as that protocol is not allowed across subnet (iebroadcast) boundaries
- B. Yes Both the monitoring instance and the application instance have to be a part of the same security group, and that security group needs to allow inbound ICMP
- C. Yes, The security group for the monitoring instance needs to allow outbound ICMP and the application instance's security group needs to allow Inbound ICMP
- D. Yes, Both the monitoring instance's security group and the application instance's security group need to allow both inbound and outbound ICMP ping packets since ICMP is not a connection-oriented protocol

Answer: D

NEW QUESTION 9

- (Topic 1)

When assessing an organization s use of AWS API access credentials which of the following three credentials should be evaluated?

Choose 3 answers

- A. Key pairs
- B. Console passwords
- C. Access keys
- D. Signing certificates
- E. Security Group memberships

Answer: ACD

Explanation:

Reference:
http://media.amazonwebservices.com/AWS_Operational_Checklists.pdf

NEW QUESTION 10

- (Topic 1)

You run a web application where web servers on EC2 Instances are In an Auto Scaling group Monitoring over the last 6 months shows that 6 web servers are necessary to handle the minimum load During the day up to 12 servers are needed Five to six days per year, the number of web servers required might go up to 15.

What would you recommend to minimize costs while being able to provide hill availability?

- A. 6 Reserved instances (heavy utilization). 6 Reserved instances (medium utilization), rest covered by On-Demand instances
- B. 6 Reserved instances (heavy utilization). 6 On-Demand instances, rest covered by Spot Instances
- C. 6 Reserved instances (heavy utilization) 6 Spot instances, rest covered by On-Demand instances
- D. 6 Reserved instances (heavy utilization) 6 Reserved instances (medium utilization) rest covered by Spot instances

Answer: B

NEW QUESTION 10

- (Topic 1)

You have a web application leveraging an Elastic Load Balancer (ELB) In front of the web servers deployed using an Auto Scaling Group Your database is running on Relational

Database Service (RDS) The application serves out technical articles and responses to them in general there are more views of an article than there are responses to the article. On occasion, an article on the site becomes extremely popular resulting in significant traffic Increases that causes the site to go down. What could you do to help alleviate the pressure on the infrastructure while maintaining availability during these events?

Choose 3 answers

- A. Leverage CloudFront for the delivery of the article
- B. Add RDS read-replicas for the read traffic going to your relational database
- C. Leverage ElastiCache for caching the most frequently used dat
- D. Use SOS to queue up the requests for the technical posts and deliver them out of the queu
- E. Use Route53 health checks to fail over to an S3 bucket for an error pag

Answer: ACE

NEW QUESTION 13

- (Topic 1)

You are running a database on an EC2 instance, with the data stored on Elastic Block Store (EBS) for persistence At times throughout the day, you are seeing large variance in the response times of the database queries Looking into the instance with the isolate command you see a lot of wait time on the disk volume that the database's data is stored on.

What two ways can you improve the performance of the database's storage while maintaining the current persistence of the data?

Choose 2 answers

- A. Move to an SSD backed instance
- B. Move the database to an EBS-Optimized Instance
- C. T Use Provisioned IOPs EBS
- D. Use the ephemeral storage on an m2 4xlarge Instance Instead

Answer: AB

NEW QUESTION 17

- (Topic 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

Explanation:

Reference:

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

NEW QUESTION 20

- (Topic 1)

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 2 answers

- A. Configure multi-factor authentication for privileged 1AM users
- B. Create 1AM users for privileged accounts
- C. Implement identity federation between your organization's Identity provider leveraging the 1AM Security Token Service
- D. Enable the 1AM single-use password policy option for privileged users

Answer: CD

NEW QUESTION 24

- (Topic 1)

When an EC2 instance that is backed by an S3-based AMI Is terminated, what happens to the data on me root volume?

- A. Data is automatically saved as an E8S volum
- B. Data is automatically saved as an ESS snapsho
- C. Data is automatically delete
- D. Data is unavailable until the instance is restarte

Answer: C

Explanation:

Reference:
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ComponentsAMIs.html>

NEW QUESTION 25

- (Topic 1)

You have an Auto Scaling group associated with an Elastic Load Balancer (ELB). You have noticed that instances launched via the Auto Scaling group are being marked unhealthy due to an ELB health check, but these unhealthy instances are not being terminated. What do you need to do to ensure that instances marked unhealthy by the ELB will be terminated and replaced?

- A. Change the thresholds set on the Auto Scaling group health check
- B. Add an Elastic Load Balancing health check to your Auto Scaling group
- C. Increase the value for the Health check interval set on the Elastic Load Balancer
- D. Change the health check set on the Elastic Load Balancer to use TCP rather than HTTP checks

Answer: B

Explanation:

Reference:
<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-add-elb-healthcheck.html>

Add an Elastic Load Balancing Health Check to your Auto Scaling Group

By default, an Auto Scaling group periodically reviews the results of EC2 instance status to determine the health state of each instance. However, if you have associated your Auto Scaling group with an Elastic Load Balancing load balancer, you can choose to use the Elastic Load Balancing health check. In this case, Auto Scaling determines the health status of your instances by checking the results of both the EC2 instance status check and the Elastic Load Balancing instance health check.

For information about EC2 instance status checks, see *Monitor Instances With Status Checks* in the *Amazon EC2 User Guide for Linux Instances*. For information about Elastic Load Balancing health checks, see *Health Check* in the *Elastic Load Balancing Developer Guide*.

This topic shows you how to add an Elastic Load Balancing health check to your Auto Scaling group, assuming that you have created a load balancer and have registered the load balancer with your Auto Scaling group. If you have not registered the load balancer with your Auto Scaling group, see *Set Up a Scaled and Load-Balanced Application*.

Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action `DescribeInstanceStatus` return any state other than `running`, the system status shows `impaired`, or the calls to Elastic Load Balancing action `DescribeInstanceHealth` returns `OutOfService` in the instance state field.

If there are multiple load balancers associated with your Auto Scaling group, Auto Scaling checks the health state of your EC2 instances by making health check calls to each load balancer. For each call, if the Elastic Load Balancing action returns any state other than `InService`, the instance is marked as unhealthy. After Auto Scaling marks an instance as unhealthy, it remains in that state, even if subsequent calls from other load balancers return an `InService` state for the same instance.

NEW QUESTION 30

- (Topic 1)

You receive a frantic call from a new DBA who accidentally dropped a table containing all your customers. Which Amazon RDS feature will allow you to reliably restore your database to within 5 minutes of when the mistake was made?

- A. Multi-AZ RDS
- B. RDS snapshots
- C. RDS read replicas
- D. RDS automated backup

Answer: D

Explanation:

Reference:
<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.BackingUpAndRestoringAmazonRDSInstances.html>

NEW QUESTION 32

- (Topic 1)

You have been asked to automate many routine systems administrator backup and recovery activities. Your current plan is to leverage AWS-managed solutions as much as possible and automate the rest with the AWS CLI and scripts. Which task would be best accomplished with a script?

- A. Creating daily EBS snapshots with a monthly rotation of snapshots
- B. Creating daily RDS snapshots with a monthly rotation of snapshots
- C. Automatically detect and stop unused or underutilized EC2 instances
- D. Automatically add Auto Scaled EC2 instances to an Amazon Elastic Load Balancer

Answer: A

NEW QUESTION 33

- (Topic 2)

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.

NEW QUESTION 38

- (Topic 2)

A user is trying to understand the ACL and policy for an S3 bucket. Which of the below mentioned policy permissions is equivalent to the WRITE ACL on a bucket?

- A. s3:GetObjectAcl
- B. s3:GetObjectVersion
- C. s3:ListBucketVersions
- D. s3:DeleteObject

Answer: D

Explanation:

Amazon S3 provides a set of operations to work with the Amazon S3 resources. Each AWS S3 bucket can have an ACL (Access Control List) or bucket policy associated with it. The WRITE ACL list allows the other AWS accounts to write/modify to that bucket. The equivalent S3 bucket policy permission for it is s3:DeleteObject.

NEW QUESTION 41

- (Topic 2)

An application is generating a log file every 5 minutes. The log file is not critical but may be required only for verification in case of some major issue. The file should be accessible over the internet whenever required. Which of the below mentioned options is a best possible storage solution for it?

- A. AWS S3
- B. AWS Glacier
- C. AWS RDS
- D. AWS RRS

Answer: D

Explanation:

Amazon S3 stores objects according to their storage class. There are three major storage classes: Standard, Reduced Redundancy Storage and Glacier. Standard is for AWS S3 and provides very high durability. However, the costs are a little higher. Glacier is for archival and the files are not available over the internet. Reduced Redundancy Storage is for less critical files. Reduced Redundancy is little cheaper as it provides less durability in comparison to S3. In this case since the log files are not mission critical files, RRS will be a better option.

NEW QUESTION 42

- (Topic 2)

A user is planning to use AWS CloudFormation. Which of the below mentioned functionalities does not help him to correctly understand CloudFormation?

- A. CloudFormation follows the DevOps model for the creation of Dev & Test
- B. AWS CloudFormation does not charge the user for its service but only charges for the AWS resources created with it
- C. CloudFormation works with a wide variety of AWS services, such as EC2, EBS, VPC, IAM, S3, RDS, ELB, etc
- D. CloudFormation provides a set of application bootstrapping scripts which enables the user to install Software

Answer: A

Explanation:

AWS CloudFormation is an application management tool which provides application modelling, deployment, configuration, management and related activities. It supports a wide variety of AWS services, such as EC2, EBS, AS, ELB, RDS, VPC, etc. It also provides application bootstrapping scripts which enable the user to install software packages or create folders. It is free of the cost and only charges the user for the services created with it. The only challenge is that it does not follow any model, such as DevOps; instead customers can define templates and use them to provision and manage the AWS resources in an orderly way.

NEW QUESTION 45

- (Topic 2)

You are managing the AWS account of a big organization. The organization has more than 1000+ employees and they want to provide access to the various services to most of the employees. Which of the below mentioned options is the best possible solution in this case?

- A. The user should create a separate IAM user for each employee and provide access to them as per the policy
- B. The user should create an IAM role and attach STS with the rol
- C. The user should attach that role to the EC2 instance and setup AWS authentication on that server
- D. The user should create IAM groups as per the organization's departments and add each user to the group for better access control
- E. Attach an IAM role with the organization's authentication service to authorize each user for various AWS services

Answer: D

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The user is managing an AWS account for an organization that already has an identity system, such as the login system for the corporate network (SSO). In this case, instead of creating individual IAM users or groups for each user who need AWS access, it may be more practical to use a proxy server to translate the user identities from

the organization network into the temporary AWS security credentials. This proxy server will attach an IAM role to the user after authentication.

NEW QUESTION 50

- (Topic 2)

A user has created a VPC with CIDR 20.0.0.0/16 with only a private subnet and VPN connection using the VPC wizard. The user wants to connect to the instance in a private subnet over SSH. How should the user define the security rule for SSH?

- A. Allow Inbound traffic on port 22 from the user's network
- B. The user has to create an instance in EC2 Classic with an elastic IP and configure the security group of a private subnet to allow SSH from that elastic IP
- C. The user can connect to a instance in a private subnet using the NAT instance
- D. Allow Inbound traffic on port 80 and 22 to allow the user to connect to a private subnet over the Internet

Answer: A

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, the user can setup a case with a VPN only subnet (private. which uses VPN access to connect with his data centre. When the user has configured this setup with Wizard, all network connections to the instances in the subnet will come from his data centre. The user has to configure the security group of the private subnet which allows the inbound traffic on SSH (port 22. from the data centre's network range.

NEW QUESTION 54

- (Topic 2)

A system admin is managing buckets, objects and folders with AWS S3. Which of the below mentioned statements is true and should be taken in consideration by the sysadmin?

- A. The folders support only ACL
- B. Both the object and bucket can have an Access Policy but folder cannot have policy
- C. Folders can have a policy
- D. Both the object and bucket can have ACL but folders cannot have ACL

Answer: A

Explanation:

A sysadmin can grant permission to the S3 objects or the buckets to any user or make objects public using the bucket policy and user policy. Both use the JSON-based access policy language. Generally if user is defining the ACL on the bucket, the objects in the bucket do not inherit it and vice a versa. The bucket policy can be defined at the bucket level which allows the objects as well as the bucket to be public with a single policy applied to that bucket. It cannot be applied at the object level. The folders are similar to objects with no content. Thus, folders can have only ACL and cannot have a policy.

NEW QUESTION 57

- (Topic 2)

A user has created an ELB with the availability zone US-East-1A. The user wants to add more zones to ELB to achieve High Availability. How can the user add more zones to the existing ELB?

- A. It is not possible to add more zones to the existing ELB
- B. The only option is to launch instances in different zones and add to ELB
- C. The user should stop the ELB and add zones and instances as required
- D. The user can add zones on the fly from the AWS console

Answer: D

Explanation:

The user has created an Elastic Load Balancer with the availability zone and wants to add more zones to the existing ELB. The user can do so in two ways:
From the console or CLI, add new zones to ELB;
Launch instances in a separate AZ and add instances to the existing ELB.

NEW QUESTION 58

- (Topic 2)

An organization is planning to create 5 different AWS accounts considering various security requirements. The organization wants to use a single payee account by using the consolidated billing option. Which of the below mentioned statements is true with respect to the above information?

- A. Master (Payee
- B. account will get only the total bill and cannot see the cost incurred by each account
- C. Master (Payee
- D. account can view only the AWS billing details of the linked accounts
- E. It is not recommended to use consolidated billing since the payee account will have access to the linked accounts
- F. Each AWS account needs to create an AWS billing policy to provide permission to the payee account

Answer: B

Explanation:

AWS consolidated billing enables the organization to consolidate payments for multiple Amazon Web Services (AWS. accounts within a single organization by making a single paying account. Consolidated billing enables the organization to see a combined view of the AWS charges incurred by each account as well as obtain a detailed cost report for each of the individual AWS accounts associated with the paying account. The payee account will not have any other access than billing data of linked accounts.

NEW QUESTION 61

- (Topic 2)

An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

- A. The organization should create a separate login ID but give the IAM users the same alias so that each one can login with their alias
- B. The organization should create each user in a separate region so that they have their own URL to login
- C. It is not possible to have the same login ID for multiple IAM users of the same account
- D. The organization should create various groups and add each user with the same login ID to different group
- E. The user can login with their own group ID

Answer: C

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Whenever the organization is creating an IAM user, there should be a unique ID for each user. It is not possible to have the same login ID for multiple users. The names of users, groups, roles, instance profiles must be alphanumeric, including the following common characters: plus (+), equal (=), comma (,), period (.), at (@), and dash (-).

NEW QUESTION 62

- (Topic 2)

A user has created an ELB with three instances. How many security groups will ELB create by default?

- A. 3
- B. 5
- C. 2
- D. 1

Answer: C

Explanation:

Elastic Load Balancing provides a special Amazon EC2 source security group that the user can use to ensure that back-end EC2 instances receive traffic only from Elastic Load Balancing. This feature needs two security groups: the source security group and a security group that defines the ingress rules for the back-end instances. To ensure that traffic only flows between the load balancer and the back-end instances, the user can add or modify a rule to the back-end security group which can limit the ingress traffic. Thus, it can come only from the source security group provided by Elastic load Balancing.

NEW QUESTION 64

- (Topic 2)

A user has created a 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 Glacier
- B. AWS Elastic Transcoder
- C. AWS Simple Notification Service
- D. AWS Simple Queue Service

Answer: D

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.

NEW QUESTION 65

- (Topic 2)

A user is planning to evaluate AWS for their internal use. The user does not want to incur any charge on his account during the evaluation. Which of the below mentioned AWS services would incur a charge if used?

- A. AWS S3 with 1 GB of storage
- B. AWS micro instance running 24 hours daily
- C. AWS ELB running 24 hours a day
- D. AWS PIOPS volume of 10 GB size

Answer: D

Explanation:

AWS is introducing a free usage tier for one year to help the new AWS customers get started in Cloud. The free tier can be used for anything that the user wants to run in the Cloud. AWS offers a handful of AWS services as a part of this which includes 750 hours of free micro instances and 750 hours of ELB. It includes the AWS S3 of 5 GB and AWS EBS general purpose volume upto 30 GB. PIOPS is not part of free usage tier.

NEW QUESTION 69

- (Topic 2)

A user has launched 10 instances from the same AMI ID using Auto Scaling. The user is trying to see the average CPU utilization across all instances of the last 2 weeks under the CloudWatch console. How can the user achieve this?

- A. View the Auto Scaling CPU metrics
- B. Aggregate the data over the instance AMI ID
- C. The user has to use the CloudWatch analyser to find the average data across instances
- D. It is not possible to see the average CPU utilization of the same AMI ID since the instance ID is different

Answer: B

Explanation:

Amazon CloudWatch is basically a metrics repository. Either the user can send the custom data or an AWS product can put metrics into the repository, and the user can retrieve the statistics based on those metrics. The statistics are metric data aggregations over specified periods of time. Aggregations are made using the namespace, metric name, dimensions, and the data point unit of measure, within the time period that is specified by the user. To aggregate the data across instances launched with AMI, the user should select the AMI ID under EC2 metrics and select the aggregate average to view the data.

NEW QUESTION 70

- (Topic 2)

A user has launched an EC2 instance from an instance store backed AMI. The infrastructure team wants to create an AMI from the running instance. Which of the below mentioned steps will not be performed while creating the AMI?

- A. Define the AMI launch permissions
- B. Upload the bundled volume
- C. Register the AMI
- D. Bundle the volume

Answer: A

Explanation:

When the user has launched an EC2 instance from an instance store backed AMI, it will need to follow certain steps, such as "Bundling the root volume", "Uploading the bundled volume" and "Register the AMI". Once the AMI is created the user can setup the launch permission. However, it is not required to setup during the launch.

NEW QUESTION 74

- (Topic 2)

A root AWS account owner is trying to understand various options to set the permission to AWS S3. Which of the below mentioned options is not the right option to grant permission for S3?

- A. User Access Policy
- B. S3 Object Access Policy
- C. S3 Bucket Access Policy
- D. S3 ACL

Answer: B

Explanation:

Amazon S3 provides a set of operations to work with the Amazon S3 resources. Managing S3 resource access refers to granting others permissions to work with S3. There are three ways the root account owner can define access with S3: S3 ACL: The user can use ACLs to grant basic read/write permissions to other AWS accounts. S3 Bucket Policy: The policy is used to grant other AWS accounts or IAM users permissions for the bucket and the objects in it. User Access Policy: Define an IAM user and assign him the IAM policy which grants him access to S3.

NEW QUESTION 76

- (Topic 2)

A user is running one instance for only 3 hours every day. The user wants to save some cost with the instance. Which of the below mentioned Reserved Instance categories is advised in this case?

- A. The user should not use RI; instead only go with the on-demand pricing
- B. The user should use the AWS high utilized RI
- C. The user should use the AWS medium utilized RI
- D. The user should use the AWS low utilized RI

Answer: A

Explanation:

The AWS Reserved Instance provides the user with an option to save some money by paying a one-time fixed amount and then save on the hourly rate. It is advisable that if the user is having 30% or more usage of an instance per day, he should go for a RI. If the user is going to use an EC2 instance for more than 2200-2500 hours per year, RI will help the user save some cost. Here, the instance is not going to run for less than 1500 hours. Thus, it is advisable that the user should use the on-demand pricing.

NEW QUESTION 81

- (Topic 2)

A user has configured Elastic Load Balancing by enabling a Secure Socket Layer (SSL) negotiation configuration known as a Security Policy. Which of the below mentioned options is not part of this secure policy while negotiating the SSL connection between the user and the client?

- A. SSL Protocols
- B. Client Order Preference
- C. SSL Ciphers

D. Server Order Preference

Answer: B

Explanation:

Elastic Load Balancing uses a Secure Socket Layer (SSL) negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. A security policy is a combination of SSL Protocols, SSL Ciphers, and the Server Order Preference option.

NEW QUESTION 83

- (Topic 2)

A user is checking the CloudWatch metrics from the AWS console. The user notices that the CloudWatch data is coming in UTC. The user wants to convert the data to a local time zone. How can the user perform this?

- A. In the CloudWatch dashboard the user should set the local timezone so that CloudWatch shows the data only in the local time zone
- B. In the CloudWatch console select the local timezone under the Time Range tab to view the data as per the local timezone
- C. The CloudWatch data is always in UTC; the user has to manually convert the data
- D. The user should have send the local timezone while uploading the data so that CloudWatch will show the data only in the local timezone

Answer: B

Explanation:

If the user is viewing the data inside the CloudWatch console, the console provides options to filter values either using the relative period, such as days/hours or using the Absolute tab where the user can provide data with a specific date and time. The console also provides the option to search using the local timezone under the time range caption in the console because the time range tab allows the user to change the time zone.

NEW QUESTION 86

- (Topic 2)

A customer is using AWS for Dev and Test. The customer wants to setup the Dev environment with Cloudformation. Which of the below mentioned steps are not required while using Cloudformation?

- A. Create a stack
- B. Configure a service
- C. Create and upload the template
- D. Provide the parameters configured as part of the template

Answer: B

Explanation:

AWS Cloudformation is an application management tool which provides application modelling, deployment, configuration, management and related activities. AWS CloudFormation introduces two concepts: the template and the stack. The template is a JSON-format, text-based file that describes all the AWS resources required to deploy and run an application. The stack is a collection of AWS resources which are created and managed as a single unit when AWS CloudFormation instantiates a template. While creating a stack, the user uploads the template and provides the data for the parameters if required.

NEW QUESTION 87

- (Topic 2)

An organization is setting up programmatic billing access for their AWS account. Which of the below mentioned services is not required or enabled when the organization wants to use programmatic access?

- A. Programmatic access
- B. AWS bucket to hold the billing report
- C. AWS billing alerts
- D. Monthly Billing report

Answer: C

Explanation:

AWS provides an option to have programmatic access to billing. Programmatic Billing Access leverages the existing Amazon Simple Storage Service (Amazon S3) APIs. Thus, the user can build applications that reference his billing data from a CSV (comma-separated value) file stored in an Amazon S3 bucket. To enable programmatic access, the user has to first enable the monthly billing report. Then the user needs to provide an AWS bucket name where the billing CSV will be uploaded. The user should also enable the Programmatic access option.

NEW QUESTION 90

- (Topic 2)

A root account owner has created an S3 bucket testmycloud. The account owner wants to allow everyone to upload the objects as well as enforce that the person who uploaded the object should manage the permission of those objects. Which is the easiest way to achieve this?

- A. The root account owner should create a bucket policy which allows the IAM users to upload the object
- B. The root account owner should create the bucket policy which allows the other account owners to set the object policy of that bucket
- C. The root account should use ACL with the bucket to allow everyone to upload the object
- D. The root account should create the IAM users and provide them the permission to upload content to the bucket

Answer: C

Explanation:

Each AWS S3 bucket and object has an ACL (Access Control List) associated with it. An ACL is a list of grants identifying the grantee and the permission granted. The user can use ACLs to grant basic read/write permissions to other AWS accounts. ACLs use an Amazon S3-specific XML schema. The user cannot grant permissions to other users in his account. ACLs are suitable for specific scenarios. For example, if a bucket owner allows other AWS accounts to upload objects, permissions to these objects can only be managed using the object ACL by the AWS account that owns the object.

NEW QUESTION 94

- (Topic 2)

A user has configured ELB with two EBS backed EC2 instances. The user is trying to understand the DNS access and IP support for ELB. Which of the below mentioned statements may not help the user understand the IP mechanism supported by ELB?

- A. The client can connect over IPV4 or IPV6 using Dualstack
- B. ELB DNS supports both IPV4 and IPV6
- C. Communication between the load balancer and back-end instances is always through IPV4
- D. The ELB supports either IPV4 or IPV6 but not both

Answer: D

Explanation:

Elastic Load Balancing supports both Internet Protocol version 6 (IPv6) and Internet Protocol version 4 (IPv4). Clients can connect to the user's load balancer using either IPv4 or IPv6 (in EC2-Classic DNS). However, communication between the load balancer and its back-end instances uses only IPv4. The user can use the Dualstack-prefixed DNS name to enable IPv6 support for communications between the client and the load balancers. Thus, the clients are able to access the load balancer using either IPv4 or IPv6 as their individual connectivity needs dictate.

NEW QUESTION 97

- (Topic 2)

An organization has added 3 of his AWS accounts to consolidated billing. One of the AWS accounts has purchased a Reserved Instance (RI) of a small instance size in the US-East-1a zone. All other AWS accounts are running instances of a small size in the same zone. What will happen in this case for the RI pricing?

- A. Only the account that has purchased the RI will get the advantage of RI pricing
- B. One instance of a small size and running in the US-East-1a zone of each AWS account will get the benefit of RI pricing
- C. Any single instance from all the three accounts can get the benefit of AWS RI pricing if they are running in the same zone and are of the same size
- D. If there are more than one instances of a small size running across multiple accounts in the same zone no one will get the benefit of RI

Answer: C

Explanation:

AWS consolidated billing enables the organization to consolidate payments for multiple Amazon Web Services (AWS) accounts within a single organization by making a single paying account. For billing purposes, consolidated billing treats all the accounts on the consolidated bill as one account. This means that all accounts on a consolidated bill can receive the hourly cost benefit of the Amazon EC2 Reserved Instances purchased by any other account. In this case only one Reserved Instance has been purchased by one account. Thus, only a single instance from any of the accounts will get the advantage of RI. AWS will implement the blended rate for each instance if more than one instance is running concurrently.

NEW QUESTION 102

- (Topic 2)

A user is trying to save some cost on the AWS services. Which of the below mentioned options will not help him save cost?

- A. Delete the unutilized EBS volumes once the instance is terminated
- B. Delete the AutoScaling launch configuration after the instances are terminated
- C. Release the elastic IP if not required once the instance is terminated
- D. Delete the AWS ELB after the instances are terminated

Answer: B

Explanation:

AWS bills the user on a pay as you go model. AWS will charge the user once the AWS resource is allocated. Even though the user is not using the resource, AWS will charge if it is in service or allocated. Thus, it is advised that once the user's work is completed he should: Terminate the EC2 instance Delete the EBS volumes Release the unutilized Elastic IPs Delete ELB The AutoScaling launch configuration does not cost the user. Thus, it will not make any difference to the cost whether it is deleted or not.

NEW QUESTION 103

- (Topic 2)

A user is trying to configure the CloudWatch billing alarm. Which of the below mentioned steps should be performed by the user for the first time alarm creation in the AWS Account Management section?

- A. Enable Receiving Billing Reports
- B. Enable Receiving Billing Alerts
- C. Enable AWS billing utility
- D. Enable CloudWatch Billing Threshold

Answer: B

Explanation:

AWS CloudWatch supports enabling the billing alarm on the total AWS charges. Before the user can create an alarm on the estimated charges, he must enable

monitoring of the estimated AWS charges, by selecting the option "Enable receiving billing alerts". It takes about 15 minutes before the user can view the billing data. The user can then create the alarms.

NEW QUESTION 108

- (Topic 2)

A sys admin is trying to understand the Auto Scaling activities. Which of the below mentioned processes is not performed by Auto Scaling?

- A. Reboot Instance
- B. Schedule Actions
- C. Replace Unhealthy
- D. Availability Zone Balancing

Answer: A

Explanation:

There are two primary types of Auto Scaling processes: Launch and Terminate, which launch or terminate instances, respectively. Some other actions performed by Auto Scaling are:

AddToLoadbalancer,
AlarmNotification, HealthCheck, AZRebalance, ReplaceUnHealthy, and ScheduledActions.

NEW QUESTION 109

- (Topic 2)

A user wants to make so that whenever the CPU utilization of the AWS EC2 instance is above 90%, the redlight of his bedroom turns on. Which of the below mentioned AWS services is helpful for this purpose?

- A. AWS CloudWatch + AWS SES
- B. AWS CloudWatch + AWS SNS
- C. Non
- D. It is not possible to configure the light with the AWS infrastructure services
- E. AWS CloudWatch and a dedicated software turning on the light

Answer: B

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and fully managed push messaging service. Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. The user can configure some sensor devices at his home which receives data on the HTTP end point (REST calls) and turn on the red light. The user can configure the CloudWatch alarm to send a notification to the AWS SNS HTTP end point (the sensor device) and it will turn the light red when there is an alarm condition.

NEW QUESTION 113

- (Topic 2)

An organization (Account ID 123412341234) has attached the below mentioned IAM policy to a user. What does this policy statement entitle the user to perform?

```
"Statement": [
{
  "Sid": "AllowUsersAllActionsForCredentials",
  "Effect": "Allow",
  "Action": [
    "iam:*AccessKey*",
  ],
  "Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
}
]
```

- A. 0
- B. 0
- C. 0
- D. 0

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 organization (Account ID 123412341234) wants some of their users to manage keys (access and secret access keys) of all IAM users, the organization should set the below mentioned policy which entitles the IAM user to modify keys of all IAM users with CLI, SDK or API.

```
"Statement": [
{
  "Sid": "AllowUsersAllActionsForCredentials",
  "Effect": "Allow",
  "Action": [
    "iam:*AccessKey*",
  ],
  "Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
}
]
```

NEW QUESTION 114

- (Topic 2)

A user has created numerous EBS volumes. What is the general limit for each AWS account for the maximum number of EBS volumes that can be created?

- A. 10000
- B. 5000
- C. 100
- D. 1000

Answer: B

Explanation:

A user can attach multiple EBS volumes to the same instance within the limits specified by his AWS account. Each AWS account has a limit on the number of Amazon EBS volumes that the user can create, and the total storage available. The default limit for the maximum number of volumes that can be created is 5000.

NEW QUESTION 115

- (Topic 2)

An organization has configured the custom metric upload with CloudWatch. The organization has given permission to its employees to upload data using CLI as well SDK. How can the user track the calls made to CloudWatch?

- A. The user can enable logging with CloudWatch which logs all the activities
- B. Use CloudTrail to monitor the API calls
- C. Create an IAM user and allow each user to log the data using the S3 bucket
- D. Enable detailed monitoring with CloudWatch

Answer: B

Explanation:

AWS CloudTrail is a web service which will allow the user to monitor the calls made to the Amazon CloudWatch API for the organization's account, including calls made by the AWS Management Console, Command Line Interface (CLI), and other services. When CloudTrail logging is turned on, CloudWatch will write log files into the Amazon S3 bucket, which is specified during the CloudTrail configuration.

NEW QUESTION 117

- (Topic 2)

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. Do not select the AZ; instead let AWS select the AZ
- C. The user can never select the availability zone while launching an instance
- D. Always select the AZ while launching an instance

Answer: B

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.

NEW QUESTION 119

- (Topic 2)

An organization has created 50 IAM users. The organization has introduced a new policy which will change the access of an IAM user. How can the organization implement this effectively so that there is no need to apply the policy at the individual user level?

- A. Use the IAM groups and add users as per their role to different groups and apply policy to group
- B. The user can create a policy and apply it to multiple users in a single go with the AWS CLI
- C. Add each user to the IAM role as per their organization role to achieve effective policy setup
- D. Use the IAM role and implement access at the role level

Answer: A

Explanation:

With AWS IAM, a group is a collection of IAM users. A group allows the user to specify permissions for a collection of users, which can make it easier to manage the permissions for those users. A group helps an organization manage access in a better way; instead of applying at the individual level, the organization can apply at the group level which is applicable to all the users who are a part of that group.

NEW QUESTION 121

- (Topic 2)

A user has setup Auto Scaling with ELB on the EC2 instances. The user wants to configure that whenever the CPU utilization is below 10%, Auto Scaling should remove one instance. How can the user configure this?

- A. The user can get an email using SNS when the CPU utilization is less than 10%. The user can use the desired capacity of Auto Scaling to remove the instance
- B. Use CloudWatch to monitor the data and Auto Scaling to remove the instances using scheduled actions
- C. Configure CloudWatch to send a notification to Auto Scaling Launch configuration when the CPU utilization is less than 10% and configure the Auto Scaling

policy to remove the instance

D. Configure CloudWatch to send a notification to the Auto Scaling group when the CPU Utilization is less than 10% and configure the Auto Scaling policy to remove the instance

Answer: D

Explanation:

Amazon CloudWatch alarms watch a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup to receive a notification on the Auto Scaling group with the CloudWatch alarm when the CPU utilization is below a certain threshold. The user can configure the Auto Scaling policy to take action for removing the instance. When the CPU utilization is below 10% CloudWatch will send an alarm to the Auto Scaling group to execute the policy.

NEW QUESTION 123

- (Topic 2)

An organization has created 50 IAM users. The organization wants that each user can change their password but cannot change their access keys. How can the organization achieve this?

- A. The organization has to create a special password policy and attach it to each user
- B. The root account owner has to use CLI which forces each IAM user to change their password on first login
- C. By default each IAM user can modify their passwords
- D. The root account owner can set the policy from the IAM console under the password policy screen

Answer: D

Explanation:

With AWS IAM, organizations can use the AWS Management Console to display, create, change or delete a password policy. As a part of managing the password policy, the user can enable all users to manage their own passwords. If the user has selected the option which allows the IAM users to modify their password, he does not need to set a separate policy for the users. This option in the AWS console allows changing only the password.

NEW QUESTION 125

- (Topic 2)

A user has setup an RDS DB with Oracle. The user wants to get notifications when someone modifies the security group of that DB. How can the user configure that?

- A. It is not possible to get the notifications on a change in the security group
- B. Configure SNS to monitor security group changes
- C. Configure event notification on the DB security group
- D. Configure the CloudWatch alarm on the DB for a change in the security group

Answer: C

Explanation:

Amazon RDS uses the Amazon Simple Notification Service to provide a notification when an Amazon RDS event occurs. These events can be configured for source categories, such as DB instance, DB security group, DB snapshot and DB parameter group. If the user is subscribed to a Configuration Change category for a DB security group, he will be notified when the DB security group is changed.

NEW QUESTION 126

- (Topic 2)

An organization wants to move to Cloud. They are looking for a secure encrypted database storage option. Which of the below mentioned AWS functionalities helps them to achieve this?

- A. AWS MFA with EBS
- B. AWS EBS encryption
- C. Multi-tier encryption with Redshift
- D. AWS S3 server side storage

Answer: B

Explanation:

AWS EBS supports encryption of the volume while creating new volumes. It also supports creating volumes from existing snapshots provided the snapshots are created from encrypted volumes. The data at rest, the I/O as well as all the snapshots of EBS will be encrypted. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between the EC2 instances and EBS storage. EBS encryption is based on the AES-256 cryptographic algorithm, which is the industry standard

NEW QUESTION 130

- (Topic 2)

You are building an online store on AWS that uses SQS to process your customer orders. Your backend system needs those messages in the same sequence the customer orders have been put in. How can you achieve that?

- A. It is not possible to do this with SQS
- B. You can use sequencing information on each message
- C. You can do this with SQS but you also need to use SWF
- D. Messages will arrive in the same order by default

Answer: B

Explanation:

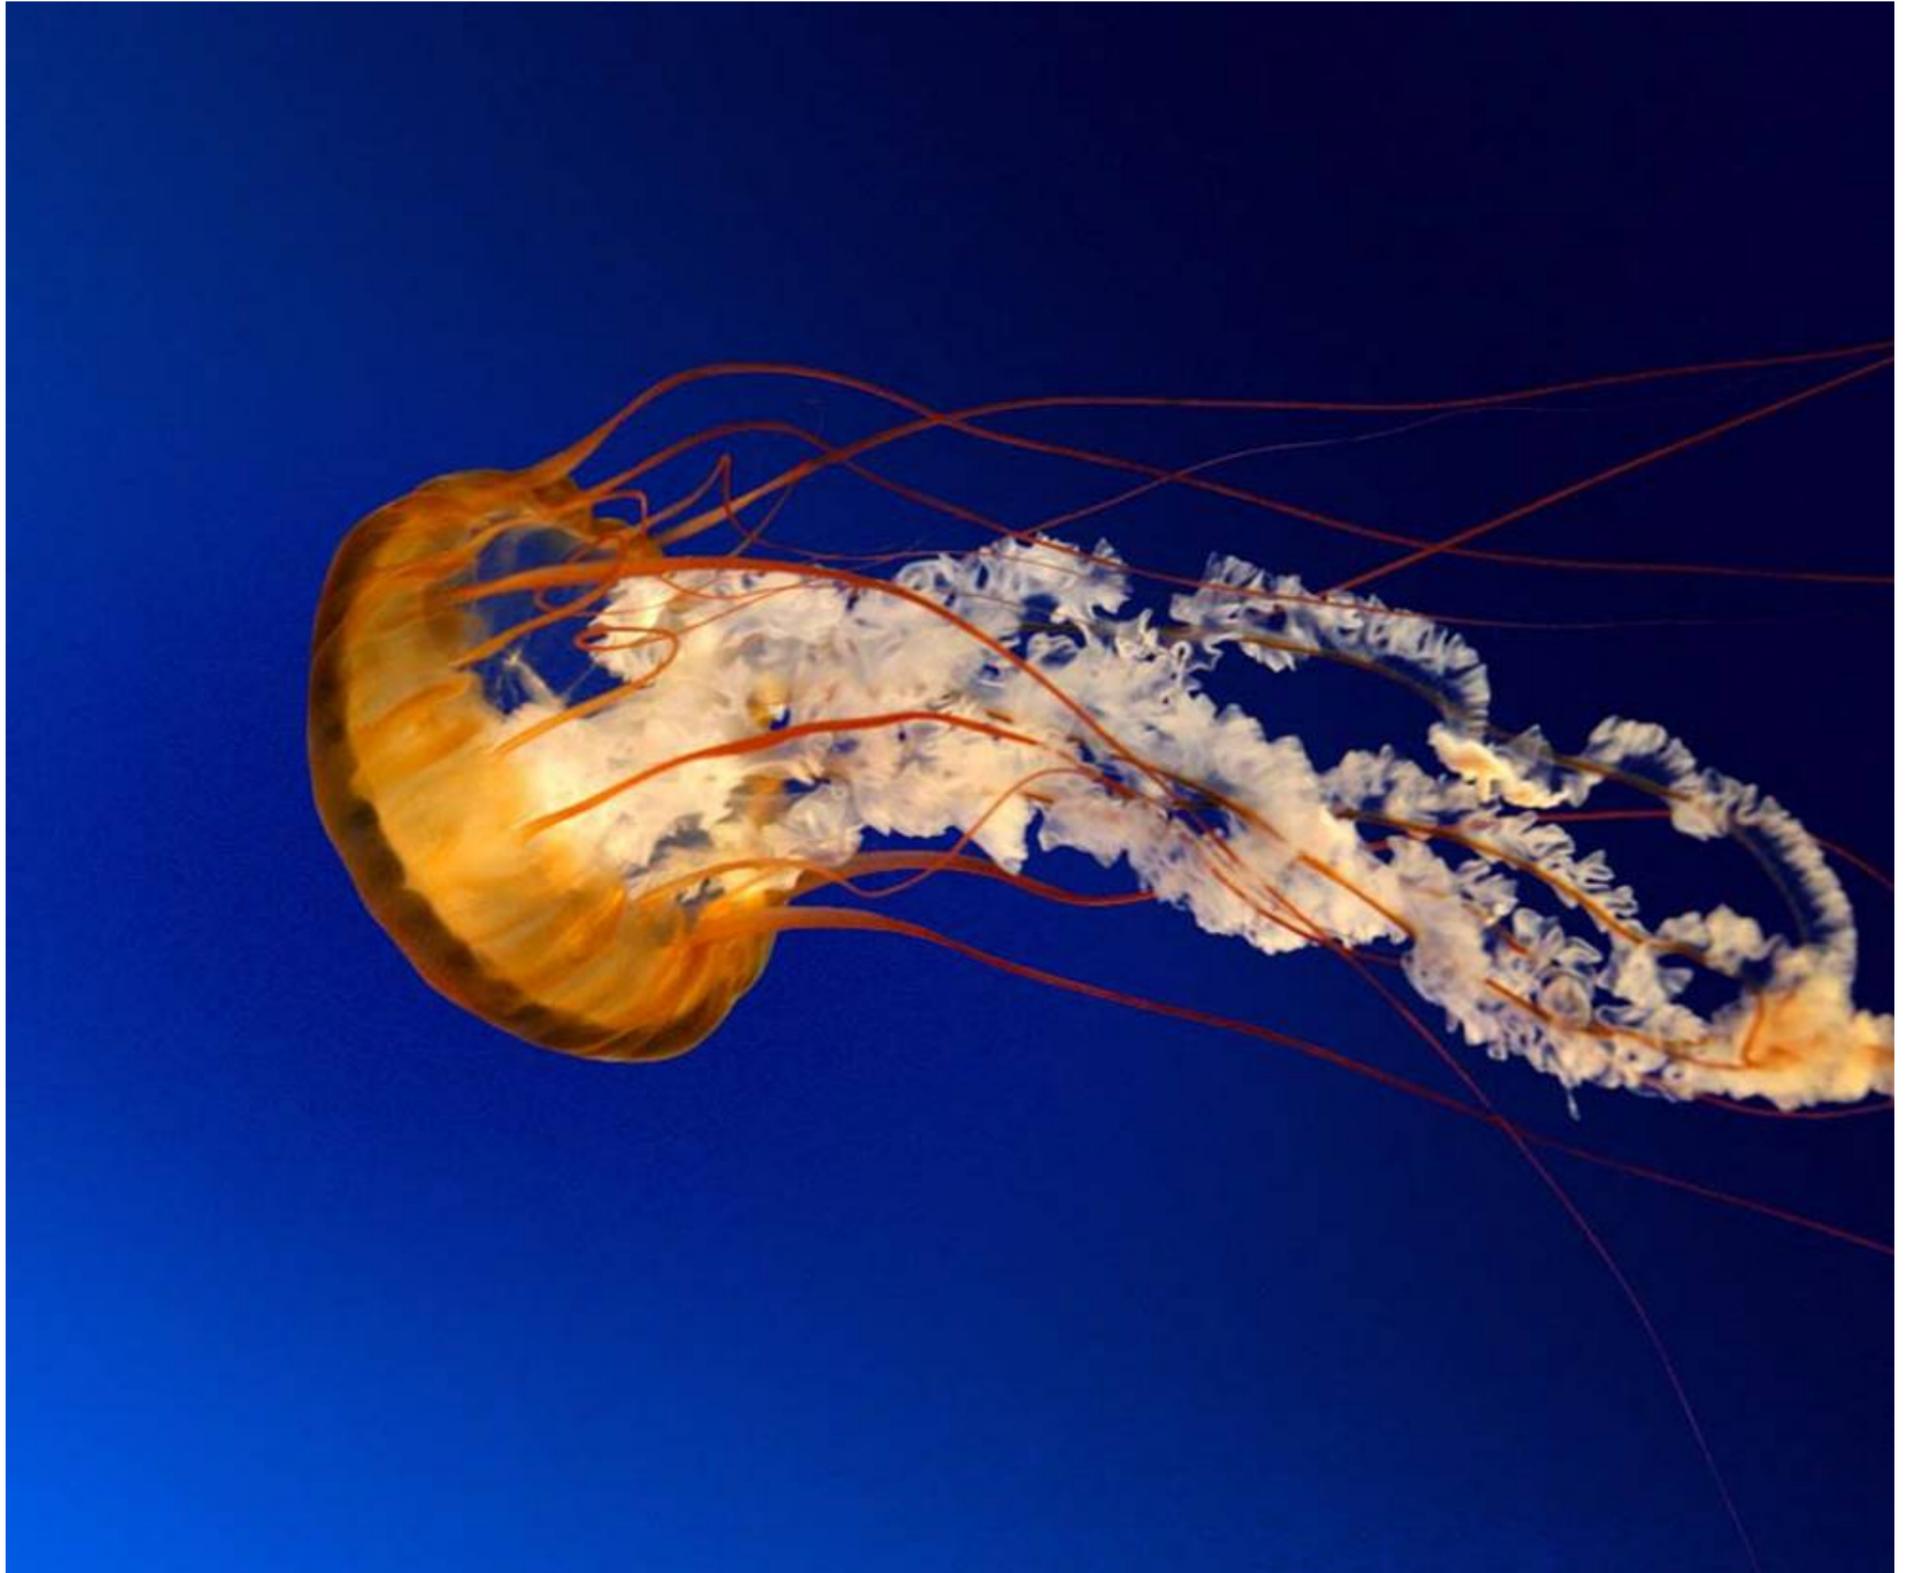
Amazon SQS is engineered to always be available and deliver messages. One of the resulting tradeoffs is that SQS does not guarantee first in, first out delivery of messages. For many distributed applications, each message can stand on its own, and as long as all messages are delivered, the order is not important. If your system requires that order be preserved, you can place sequencing information in each message, so that you can reorder the messages when the queue returns them.

NEW QUESTION 133

- (Topic 2)

A user has configured the AWS CloudWatch alarm for estimated usage charges in the US East region. Which of the below mentioned statements is not true with respect to the estimated charges?

Exhibit:



- A. It will store the estimated charges data of the last 14 days
- B. It will include the estimated charges of every AWS service
- C. The metric data will represent the data of all the regions
- D. The metric data will show data specific to that region

Answer: D

Explanation:

When the user has enabled the monitoring of estimated charges for the AWS account with AWS CloudWatch, the estimated charges are calculated and sent several times daily to CloudWatch in the form of metric data. This data will be stored for 14 days. The billing metric data is stored in the US East (Northern Virginia) Region and represents worldwide charges. This data also includes the estimated charges for every service in AWS used by the user, as well as the estimated overall AWS charges.

NEW QUESTION 137

- (Topic 2)

An organization is using cost allocation tags to find the cost distribution of different departments and projects. One of the instances has two separate tags with the key/ value as "InstanceName/HR", "CostCenter/HR". What will AWS do in this case?

- A. InstanceName is a reserved tag for AW
- B. Thus, AWS will not allow this tag
- C. AWS will not allow the tags as the value is the same for different keys
- D. AWS will allow tags but will not show correctly in the cost allocation report due to the same value of the two separate keys
- E. AWS will allow both the tags and show properly in the cost distribution report

Answer: D

Explanation:

AWS provides cost allocation tags to categorize and track the AWS costs. When the user applies tags to his AWS resources, AWS generates a cost allocation report as a comma-separated value (CSV file) with the usage and costs aggregated by those tags. Each tag will have a key-value and can be applied to services, such as EC2, S3, RDS, EMR, etc. It is required that the key should be different for each tag. The value can be the same for different keys. In this case since the value is different, AWS will properly show the distribution report with the correct values.

NEW QUESTION 141

- (Topic 2)

A sys admin is trying to understand EBS snapshots. Which of the below mentioned statements will not be useful to the admin to understand the concepts about a snapshot?

- A. The snapshot is synchronous
- B. It is recommended to stop the instance before taking a snapshot for consistent data
- C. The snapshot is incremental
- D. The snapshot captures the data that has been written to the hard disk when the snapshot command was executed

Answer: A

Explanation:

The AWS snapshot is a point in time backup of an EBS volume. When the snapshot command is executed it will capture the current state of the data that is written on the drive and take a backup. For a better and consistent snapshot of the root EBS volume, AWS recommends stopping the instance. For additional volumes it is recommended to unmount the device. The snapshots are asynchronous and incremental.

NEW QUESTION 143

- (Topic 3)

A user has deployed an application on an EBS backed EC2 instance. For a better performance of application, it requires dedicated EC2 to EBS traffic. How can the user achieve this?

- A. Launch the EC2 instance as EBS dedicated with PIOPS EBS
- B. Launch the EC2 instance as EBS enhanced with PIOPS EBS
- C. Launch the EC2 instance as EBS dedicated with PIOPS EBS
- D. Launch the EC2 instance as EBS optimized with PIOPS EBS

Answer: D

Explanation:

Any application which has performance sensitive workloads and requires minimal variability with dedicated EC2 to EBS traffic should use provisioned IOPS EBS volumes, which are attached to an EBS-optimized EC2 instance or it should use an instance with 10 Gigabit network connectivity. Launching an instance that is EBS optimized provides the user with a dedicated connection between the EC2 instance and the EBS volume.

NEW QUESTION 144

- (Topic 3)

A user is trying to understand the CloudWatch metrics for the AWS services. It is required that the user should first understand the namespace for the AWS services. Which of the below mentioned is not a valid namespace for the AWS services?

- A. AWS/StorageGateway
- B. AWS/CloudTrail
- C. AWS/ElastiCache
- D. AWS/SWF

Answer: B

Explanation:

Amazon CloudWatch is basically a metrics repository. The AWS product puts metrics into this repository, and the user can retrieve the data or statistics based on those metrics. To distinguish the data for each service, the CloudWatch metric has a namespace. Namespaces are containers for metrics. All AWS services that provide the Amazon CloudWatch data use a namespace string, beginning with "AWS/". All the services which are supported by CloudWatch will have some namespace. CloudWatch does not monitor CloudTrail. Thus, the namespace "AWS/CloudTrail" is incorrect.

NEW QUESTION 149

- (Topic 3)

A sys admin has enabled logging on ELB. Which of the below mentioned fields will not be a part of the log file name?

- A. Load Balancer IP
- B. EC2 instance IP
- C. S3 bucket name
- D. Random string

Answer: B

Explanation:

Elastic Load Balancing access logs capture detailed information for all the requests made to the load balancer. Elastic Load Balancing publishes a log file from each load balancer node at the interval that the user has specified. The load balancer can deliver multiple logs for the same period. Elastic Load Balancing creates log file names in the following format: "{Bucket}/{Prefix}/AWSLogs/{AWS AccountID}/elasticloadbalancing/{Region}/{Year}/{Month}/{Day}/{AWS Account ID}_elasticloadbalancing_{Region}_{Load Balancer Name}_{End Time}_{Load Balancer IP}_{Random String}.log"

NEW QUESTION 150

- (Topic 3)

A user is trying to connect to a running EC2 instance using SSH. However, the user gets an Unprotected Private Key File error. Which of the below mentioned options can be a possible reason for rejection?

- A. The private key file has the wrong file permission
- B. The ppk file used for SSH is read only
- C. The public key file has the wrong permission
- D. The user has provided the wrong user name for the OS login

Answer: A

Explanation:

While doing SSH to an EC2 instance, if you get an Unprotected Private Key File error it means that the private key file's permissions on your computer are too open. Ideally the private key should have the Unix permission of 0400. To fix that, run the command: `chmod 0400 /path/to/private.key`

NEW QUESTION 152

- (Topic 3)

A user has created a VPC with CIDR 20.0.0.0/24. The user has used all the IPs of CIDR and wants to increase the size of the VPC. The user has two subnets: public (20.0.0.0/28. and private (20.0.1.0/28.. How can the user change the size of the VPC?

- A. The user can delete all the instances of the subne
- B. Change the size of the subnets to 20.0.0.0/32 and 20.0.1.0/32, respective
- C. Then the user can increase the size of the VPC using CLI
- D. It is not possible to change the size of the VPC once it has been created
- E. The user can add a subnet with a higher range so that it will automatically increase the size of the VPC
- F. The user can delete the subnets first and then modify the size of the VPC

Answer: B

Explanation:

Once the user has created a VPC, he cannot change the CIDR of that VPC. The user has to terminate all the instances, delete the subnets and then delete the VPC. Create a new VPC with a higher size and launch instances with the newly created VPC and subnets.

NEW QUESTION 157

- (Topic 3)

A user has launched an EC2 instance store backed instance in the US-East-1a zone. The user created AMI #1 and copied it to the Europe region. After that, the user made a few updates to the application running in the US-East-1a zone. The user makes an AMI#2 after the changes. If the user launches a new instance in Europe from the AMI #1 copy, which of the below mentioned statements is true?

- A. The new instance will have the changes made after the AMI copy as AWS just copies the reference of the original AMI during the copyin
- B. Thus, the copied AMI will have all the updated data
- C. The new instance will have the changes made after the AMI copy since AWS keeps updating the AMI
- D. It is not possible to copy the instance store backed AMI from one region to another
- E. The new instance in the EU region will not have the changes made after the AMI copy

Answer: D

Explanation:

Within EC2, when the user copies an AMI, the new AMI is fully independent of the source AMI; there is no link to the original (source. AMI. The user can modify the source AMI without affecting the new AMI and vice a versa. Therefore, in this case even if the source AMI is modified, the copied AMI of the EU region will not have the changes. Thus, after copy the user needs to copy the new source AMI to the destination region to get those changes.

NEW QUESTION 159

- (Topic 3)

An AWS account owner has setup multiple IAM users. One IAM user only has CloudWatch access. He has setup the alarm action which stops the EC2 instances when the CPU utilization is below the threshold limit. What will happen in this case?

- A. It is not possible to stop the instance using the CloudWatch alarm
- B. CloudWatch will stop the instance when the action is executed
- C. The user cannot set an alarm on EC2 since he does not have the permission
- D. The user can setup the action but it will not be executed if the user does not have EC2 rights

Answer: D

Explanation:

Amazon CloudWatch alarms watch a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup an action which stops the instances when their CPU utilization is below a certain threshold for a certain period of time. The EC2 action can either terminate or stop the instance as part of the EC2 action. If the IAM user has read/write permissions for Amazon CloudWatch but not for Amazon EC2, he can still create an alarm. However, the stop or terminate actions will not be performed on the Amazon EC2 instance.

NEW QUESTION 162

- (Topic 3)

A user is sending the data to CloudWatch using the CloudWatch API. The user is sending data 90 minutes in the future. What will CloudWatch do in this case?

- A. CloudWatch will accept the data
- B. It is not possible to send data of the future
- C. It is not possible to send the data manually to CloudWatch
- D. The user cannot send data for more than 60 minutes in the future

Answer: A

Explanation:

With Amazon CloudWatch, each metric data point must be marked with a time stamp. The user can send the data using CLI but the time has to be in the UTC format. If the user does not provide the time, CloudWatch will take the data received time in the UTC timezone. The time stamp sent by the user can be up to two weeks in the past and up to two hours into the future.

NEW QUESTION 164

- (Topic 3)

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. Which of the below mentioned statements is true with respect to this scenario?

- A. The instance will always have a public DNS attached to the instance by default
- B. The user can directly attach an elastic IP to the instance
- C. The instance will never launch if the public IP is not assigned
- D. The user would need to create an internet gateway and then attach an elastic IP to the instance to connect from internet

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. The user cannot connect to the instance from the internet. If the user wants an elastic IP to connect to the instance from the internet he should create an internet gateway and assign an elastic IP to instance.

NEW QUESTION 166

- (Topic 3)

Which of the following statements about this S3 bucket policy is true?

```
{
  "Id": "IPAllowPolicy",
  "Statement": [
    {
      "Sid": "IPAllow",
      "Action": "s3:*",
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::mybucket/*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": "192.168.100.0/24"
        },
        "NotIpAddress": {
          "aws:SourceIp": "192.168.100.188/32"
        }
      }
    }
  ],
  "Principal": {
    "AWS": [
      "*"
    ]
  }
}
```

- A. Denies the server with the IP address 192.166 100.0 full access to the "mybucket" bucket
- B. Denies the server with the IP address 192.166 100.188 full access to the "mybucket bucket
- C. Grants all the servers within the 192 168 100 0/24 subnet full access to the "mybucket" bucket
- D. Grants all the servers within the 192 168 100 188/32 subnet full access to the "mybucket" bucket

Answer: C

NEW QUESTION 169

- (Topic 3)

A user has created an Auto Scaling group with default configurations from CLI. The user wants to setup the CloudWatch alarm on the EC2 instances, which are launched by the Auto Scaling group. The user has setup an alarm to monitor the CPU utilization every minute. Which of the below mentioned statements is true?

- A. It will fetch the data at every minute but the four data points [corresponding to 4 minutes] will not have value since the EC2 basic monitoring metrics are collected every five minutes
- B. It will fetch the data at every minute as detailed monitoring on EC2 will be enabled by the default launch configuration of Auto Scaling
- C. The alarm creation will fail since the user has not enabled detailed monitoring on the EC2 instances
- D. The user has to first enable detailed monitoring on the EC2 instances to support alarm monitoring at every minute

Answer: B

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. To enable detailed instance monitoring for a new Auto Scaling group, the user does not need to take any extra steps. When the user creates an Auto Scaling launch config using CLI, each launch configuration contains a flag named InstanceMonitoring.Enabled. The default value of this flag is true. Thus, by default detailed monitoring will be enabled for Auto Scaling as well as for all the instances launched by that Auto Scaling group.

NEW QUESTION 173

- (Topic 3)

You have a business-to-business web application running in a VPC consisting of an Elastic Load Balancer (ELB), web servers, application servers and a database. Your web application should only accept traffic from pre-defined customer IP addresses.

Which two options meet this security requirement? Choose 2 answers A. Configure web server VPC security groups to allow traffic from your customers' IPs

- A. Configure your web servers to filter traffic based on the ELB's "X-forwarded-for" header
- B. Configure ELB security groups to allow traffic from your customers' IPs and deny all outbound traffic
- C. Configure a VPC NACL to allow web traffic from your customers' IPs and deny all outbound traffic

Answer: AB

NEW QUESTION 174

- (Topic 3)

An organization has configured two single availability zones. The Auto Scaling groups are configured in separate zones. The user wants to merge the groups such that one group spans across multiple zones. How can the user configure this?

- A. Run the command `as-join-auto-scaling-group` to join the two groups
- B. Run the command `as-update-auto-scaling-group` to configure one group to span across zones and delete the other group
- C. Run the command `as-copy-auto-scaling-group` to join the two groups
- D. Run the command `as-merge-auto-scaling-group` to merge the groups

Answer: B

Explanation:

If the user has configured two separate single availability zone Auto Scaling groups and wants to merge them then he should update one of the groups and delete the other one. While updating the first group it is recommended that the user should increase the size of the minimum, maximum and desired capacity as a summation of both the groups.

NEW QUESTION 175

- (Topic 3)

Your mission is to create a lights-out datacenter environment, and you plan to use AWS OpsWorks to accomplish this. First you created a stack and added an App Server layer with an instance running in it. Next you added an application to the instance, and now you need to deploy a MySQL RDS database instance.

Which of the following answers accurately describe how to add a backend database server to an OpsWorks stack? Choose 3 answers

- A. Add a new database layer and then add recipes to the deploy actions of the database and App Server layer
- B. Use OpsWorks' "Clone Stack" feature to create a second RDS stack in another Availability Zone for redundancy in the event of a failure in the Primary A
- C. To switch to the secondary RDS instance, set the `[:database]` attributes to values that are appropriate for your server which you can do by using custom JSO
- D. The variables that characterize the RDS database connection—host, user, and so on—are set using the corresponding values from the deploy JSON's `[:deploy][:app_name][:database]` attribute
- E. Cookbook attributes are stored in a repository, so OpsWorks requires that the "password": "your_password" attribute for the RDS instance must be encrypted using at least a 256-bit ke
- F. Set up the connection between the app server and the RDS layer by using a custom recip
- G. The recipe configures the app server as required, typically by creating a configuration fil
- H. The recipe gets the connection data such as the host and database name from a set of attributes in the stack configuration and deployment JSON that AWS OpsWorks installs on every instanc

Answer: BCE

NEW QUESTION 178

- (Topic 3)

A user has launched an RDS MySQL DB with the Multi AZ feature. The user has scheduled the scaling of instance storage during maintenance window. What is the correct order of events during maintenance window?

- Perform maintenance on standby
- Promote standby to primary

Perform maintenance on original primary
Promote original master back as primary

- A. 1, 2, 3, 4
- B. 1, 2, 3
- C. 2, 3, 1, 4

Answer: B

Explanation:

Running MySQL on the RDS DB instance as a Multi-AZ deployment can help the user reduce the impact of a maintenance event, as the Amazon will conduct maintenance by following the steps in the below mentioned order: Perform maintenance on standby Promote standby to primary Perform maintenance on original primary, which becomes the new standby.

NEW QUESTION 179

- (Topic 3)

A user has created a VPC with public and private subnets using the VPC wizard. Which of the below mentioned statements is true in this scenario?

- A. The AWS VPC will automatically create a NAT instance with the micro size
- B. VPC bounds the main route table with a private subnet and a custom route table with a public subnet
- C. The user has to manually create a NAT instance
- D. VPC bounds the main route table with a public subnet and a custom route table with a private subnet

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. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance of a smaller or higher size, respectively. The VPC has an implied router and the VPC wizard updates the main route table used with the private subnet, creates a custom route table and associates it with the public subnet.

NEW QUESTION 183

- (Topic 3)

A user has setup a CloudWatch alarm on the EC2 instance for CPU utilization. The user has setup to receive a notification on email when the CPU utilization is higher than 60%. The user is running a virus scan on the same instance at a particular time. The user wants to avoid receiving an email at this time. What should the user do?

- A. Remove the alarm
- B. Disable the alarm for a while using CLI
- C. Modify the CPU utilization by removing the email alert
- D. Disable the alarm for a while using the console

Answer: B

Explanation:

Amazon CloudWatch alarm watches a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. When the user has setup an alarm and it is known that for some unavoidable event the status may change to Alarm, the user can disable the alarm using the DisableAlarmActions API or from the command line `mon-disable-alarm-actions`.

NEW QUESTION 185

- (Topic 3)

A user is using the AWS SQS to decouple the services. Which of the below mentioned operations is not supported by SQS?

- A. SendMessageBatch
- B. DeleteMessageBatch
- C. CreateQueue
- D. DeleteMessageQueue

Answer: D

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 perform the following set of operations using the Amazon SQS: CreateQueue, ListQueues, DeleteQueue, SendMessage, SendMessageBatch, ReceiveMessage, DeleteMessage, DeleteMessageBatch, ChangeMessageVisibility, ChangeMessageVisibilityBatch, SetQueueAttributes, GetQueueAttributes, GetQueueUrl, AddPermission and RemovePermission. Operations can be performed only by the AWS account owner or an AWS account that the account owner has delegated to.

NEW QUESTION 188

- (Topic 3)

A user has enabled detailed CloudWatch monitoring with the AWS Simple Notification Service. Which of the below mentioned statements helps the user

understand detailed monitoring better?

- A. SNS will send data every minute after configuration
- B. There is no need to enable since SNS provides data every minute
- C. AWS CloudWatch does not support monitoring for SNS
- D. SNS cannot provide data every minute

Answer: D

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. The AWS SNS service sends data every 5 minutes. Thus, it supports only the basic monitoring. The user cannot enable detailed monitoring with SNS.

NEW QUESTION 191

- (Topic 3)

A user is having data generated randomly based on a certain event. The user wants to upload that data to CloudWatch. It may happen that event may not have data generated for some period due to andomness. Which of the below mentioned options is a recommended option for this case?

- A. For the period when there is no data, the user should not send the data at all
- B. For the period when there is no data the user should send a blank value
- C. For the period when there is no data the user should send the value as 0
- D. The user must upload the data to CloudWatch as having no data for some period will cause an error at CloudWatch monitoring

Answer: C

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. When the user data is more random and not generated at regular intervals, there can be a period which has no associated data. The user can either publish the zero (0) Value for that period or not publish the data at all. It is recommended that the user should publish zero instead of no value to monitor the health of the application. This is helpful in an alarm as well as in the generation of the sample data count.

NEW QUESTION 195

- (Topic 3)

A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24. What will happen in this scenario?

- A. The VPC will modify the first subnet CIDR automatically to allow the second subnet IP range
- B. It is not possible to create a subnet with the same CIDR as VPC
- C. The second subnet will be created
- D. It will throw a CIDR overlaps error

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. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet.

NEW QUESTION 198

- (Topic 3)

A user wants to upload a complete folder to AWS S3 using the S3 Management console. How can the user perform this activity?

- A. Just drag and drop the folder using the flash tool provided by S3
- B. Use the Enable Enhanced Folder option from the S3 console while uploading objects
- C. The user cannot upload the whole folder in one go with the S3 management console
- D. Use the Enable Enhanced Uploader option from the S3 console while uploading objects

Answer: D

Explanation:

AWS S3 provides a console to upload objects to a bucket. The user can use the file upload screen to upload the whole folder in one go by clicking on the Enable Enhanced Uploader option. When the user uploads a folder, Amazon S3 uploads all the files and subfolders from the specified folder to the user's bucket. It then assigns a key value that is a combination of the uploaded file name and the folder name.

NEW QUESTION 203

- (Topic 3)

A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling AlarmNotification (which notifies Auto Scaling for CloudWatch alarms) process for a while. What will Auto Scaling do during this period?

- A. AWS will not receive the alarms from CloudWatch
- B. AWS will receive the alarms but will not execute the Auto Scaling policy
- C. Auto Scaling will execute the policy but it will not launch the instances until the process is resumed
- D. It is not possible to suspend the AlarmNotification process

Answer: B

Explanation:

Auto Scaling performs various processes, such as Launch, Terminate Alarm Notification etc. The user can also suspend individual process. The AlarmNotification process type accepts notifications from the Amazon CloudWatch alarms that are associated with the Auto Scaling group. If the user suspends this process type, Auto Scaling will not automatically execute the scaling policies that would be triggered by the alarms.

NEW QUESTION 207

- (Topic 3)

A user has configured Auto Scaling with 3 instances. The user had created a new AMI after updating one of the instances. If the user wants to terminate two specific instances to ensure that Auto Scaling launches an instances with the new launch configuration, which command should he run?

- A. `as-delete-instance-in-auto-scaling-group <Instance ID> --no-decrement-desired-capacity`
- B. `as-terminate-instance-in-auto-scaling-group <Instance ID> --update-desired-capacity`
- C. `as-terminate-instance-in-auto-scaling-group <Instance ID> --decrement-desired-capacity`
- D. `as-terminate-instance-in-auto-scaling-group <Instance ID> --no-decrement-desired-capacity`

Answer: D

Explanation:

The Auto Scaling command `as-terminate-instance-in-auto-scaling-group <Instance ID>` will terminate the specific instance ID. The user is required to specify the parameter `--no-decrement-desired-capacity` to ensure that it launches a new instance from the launch config after terminating the instance. If the user specifies the parameter `--decrement-desired-capacity` then Auto Scaling will terminate the instance and decrease the desired capacity by 1.

NEW QUESTION 211

- (Topic 3)

A user has created a VPC with public and private subnets using the VPC wizard. The user has not launched any instance manually and is trying to delete the VPC. What will happen in this scenario?

- A. It will not allow to delete the VPC as it has subnets with route tables
- B. It will not allow to delete the VPC since it has a running route instance
- C. It will terminate the VPC along with all the instances launched by the wizard
- D. It will not allow to delete the VPC since it has a running NAT instance

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. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance with an elastic IP. If the user is trying to delete the VPC it will not allow as the NAT instance is still running.

NEW QUESTION 215

- (Topic 3)

A user runs the command `dd if=/dev/zero of=/dev/xvdfbs=1M` on a fresh blank EBS volume attached to a Linux instance. Which of the below mentioned activities is the user performing with the command given above?

- A. Creating a file system on the EBS volume
- B. Mounting the device to the instance
- C. Pre warming the EBS volume
- D. Formatting the EBS volume

Answer: C

Explanation:

When the user creates a new EBS volume and is trying to access it for the first time it will encounter reduced IOPS due to wiping or initiating of the block storage. To avoid this as well as achieve the best performance it is required to pre warm the EBS volume. For a blank volume attached with a Linux OS, the "dd" command is used to write to all the blocks on the device. In the command `dd if=/dev/zero of=/dev/xvdfbs=1M` the parameter "if =import file" should be set to one of the Linux virtual devices, such as /dev/zero. The "of=output file" parameter should be set to the drive that the user wishes to warm. The "bs" parameter sets the block size of the write operation; for optimal performance, this should be set to 1 MB.

NEW QUESTION 218

- (Topic 3)

George has shared an EC2 AMI created in the US East region from his AWS account with Stefano. George copies the same AMI to the US West region. Can Stefano access the copied AMI of George's account from the US West region?

- A. No, copy AMI does not copy the permission
- B. It is not possible to share the AMI with a specific account
- C. Yes, since copy AMI copies all private account sharing permissions
- D. Yes, since copy AMI copies all the permissions attached with the AMI

Answer: A

Explanation:

Within EC2, when the user copies an AMI, the new AMI is fully independent of the source AMI; there is no link to the original (source) AMI. AWS does not copy launch the permissions, user-defined tags or the Amazon S3 bucket permissions from the source AMI to the new AMI. Thus, in this case by default Stefano will not have access to the AMI in the US West region.

NEW QUESTION 221

- (Topic 3)

An organization (Account ID 123412341234) has attached the below mentioned IAM policy to a user. What does this policy statement entitle the user to perform?

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowUsersAllActionsForCredentials",
    "Effect": "Allow",
    "Action": [
      "iam:*LoginProfile",
      "iam:*AccessKey*",
      "iam:*SigningCertificate*"
    ],
    "Resource": ["arn:aws:iam::123412341234:user/${aws:username}"]
  }]
}
```

- A. The policy allows the IAM user to modify all IAM user's credentials using the console, SDK, CLI or APIs
- B. The policy will give an invalid resource error
- C. The policy allows the IAM user to modify all credentials using only the console
- D. The policy allows the user to modify all IAM user's password, sign in certificates and access keys using only CLI, SDK or APIs

Answer: D

Explanation:

WS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (Account ID 123412341234) wants some of their users to manage credentials (access keys, password, and sign in certificates) of all IAM users, they should set an applicable policy to that user or group of users. The below mentioned policy allows the IAM user to modify the credentials of all IAM user's using only CLI, SDK or APIs. The user cannot use the AWS console for this activity since he does not have list permission for the IAM users.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowUsersAllActionsForCredentials",
    "Effect": "Allow"
    "Action": [
      "iam:*LoginProfile",
      "iam:*AccessKey*",
      "iam:*SigningCertificate*"
    ],
    "Resource": ["arn:aws:iam::123412341234:user/${aws:username}"]
  }]
}
Amazon AWS-SysOps : Practice Test
}}
```

NEW QUESTION 226

- (Topic 3)

The compliance department within your multi-national organization requires that all data for your customers that reside in the European Union (EU) must not leave the EU and also

data for customers that reside in the US must not leave the US without explicit authorization.

What must you do to comply with this requirement for a web based profile management application running on EC2?

- A. Run EC2 instances in multiple AWS Availability Zones in single Region and leverage an Elastic Load Balancer with session stickiness to route traffic to the appropriate zone to create their profile
- B. Run EC2 instances in multiple Regions and leverage Route 53's Latency Based Routing capabilities to route traffic to the appropriate region to create their profile
- C. Run EC2 instances in multiple Regions and leverage a third party data provider to determine if a user needs to be redirect to the appropriate region to create their profile
- D. Run EC2 instances in multiple AWS Availability Zones in a single Region and leverage a third party data provider to determine if a user needs to be redirect to the appropriate zone to create their profile

Answer: C

NEW QUESTION 231

- (Topic 3)

A user has launched an EBS backed EC2 instance in the US-East-1a region. The user stopped the instance and started it back after 20 days. AWS throws up an 'InsufficientInstanceCapacity' error. What can be the possible reason for this?

- A. AWS does not have sufficient capacity in that availability zone
- B. AWS zone mapping is changed for that user account
- C. There is some issue with the host capacity on which the instance is launched
- D. The user account has reached the maximum EC2 instance limit

Answer: A

Explanation:

When the user gets an 'InsufficientInstanceCapacity' error while launching or starting an EC2 instance, it means that AWS does not currently have enough available capacity to service the user request. If the user is requesting a large number of instances, there might not be enough server capacity to host them. The user can either try again later, by specifying a smaller number of instances or changing the availability zone if launching a fresh instance.

NEW QUESTION 235

- (Topic 3)

A user is using CloudFormation to launch an EC2 instance and then configure an application after the instance is launched. The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly. How can the user configure this?

- A. It is not possible that the stack creation will wait until one service is created and launched
- B. The user can use the HoldCondition resource to wait for the creation of the other dependent resources
- C. The user can use the DependentCondition resource to hold the creation of the other dependent resources
- D. The user can use the WaitCondition resource to hold the creation of the other dependent resources

Answer: D

Explanation:

AWS CloudFormation is an application management tool which provides application modelling, deployment, configuration, management and related activities. AWS CloudFormation provides a WaitCondition resource which acts as a barrier and blocks the creation of other resources until a completion signal is received from an external source, such as a user application or management system.

NEW QUESTION 237

- (Topic 3)

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

Answer: B

Explanation:

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.

NEW QUESTION 241

- (Topic 3)

A user is creating a CloudFormation stack. Which of the below mentioned limitations does not hold true for CloudFormation?

- A. One account by default is limited to 100 templates
- B. The user can use 60 parameters and 60 outputs in a single template
- C. The template, parameter, output, and resource description fields are limited to 4096 characters
- D. One account by default is limited to 20 stacks

Answer: A

Explanation:

AWS CloudFormation is an application management tool which provides application modelling, deployment, configuration, management and related activities. The limitations given below apply to the CloudFormation template and stack. There are no limits to the number of templates but each AWS CloudFormation account is limited to a maximum of 20 stacks by default. The Template, Parameter, Output, and Resource description fields are limited to 4096 characters. The user can include up to 60 parameters and 60 outputs in a template.

NEW QUESTION 244

- (Topic 3)

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

Answer: D

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.

NEW QUESTION 249

- (Topic 3)

An organization is planning to create a user with IAM. They are trying to understand the limitations of IAM so that they can plan accordingly. Which of the below mentioned statements is not true with respect to the limitations of IAM?

- A. One IAM user can be a part of a maximum of 5 groups
- B. The organization can create 100 groups per AWS account
- C. One AWS account can have a maximum of 5000 IAM users
- D. One AWS account can have 250 roles

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. The default maximums for each of the IAM entities is given below: Groups per AWS account: 100 Users per AWS account: 5000 Roles per AWS account: 250 Number of groups per user: 10 (that is, one user can be part of these many groups).

NEW QUESTION 253

- (Topic 3)

Your business is building a new application that will store its entire customer database on a RDS MySQL database, and will have various applications and users that will query that data for different purposes.

Large analytics jobs on the database are likely to cause other applications to not be able to get the query results they need to, before time out. Also, as your data grows, these analytics jobs will start to take more time, increasing the negative effect on the other applications.

How do you solve the contention issues between these different workloads on the same data?

- A. Enable Multi-AZ mode on the RDS instance
- B. Use ElastiCache to offload the analytics job data
- C. Create RDS Read-Replicas for the analytics work
- D. Run the RDS instance on the largest size possible

Answer: B

NEW QUESTION 255

- (Topic 3)

A user has launched an EC2 instance from an instance store backed AMI. If the user restarts the instance, what will happen to the ephemeral storage data?

- A. All the data will be erased but the ephemeral storage will stay connected
- B. All data will be erased and the ephemeral storage is released
- C. It is not possible to restart an instance launched from an instance store backed AMI
- D. The data is preserved

Answer: D

Explanation:

A user can reboot an EC2 instance using the AWS console, the Amazon EC2 CLI or the Amazon EC2 API. Rebooting an instance is equivalent to rebooting an operating system. However, it is recommended that the user use Amazon EC2 to reboot the instance instead of running the operating system reboot command from the instance. When an instance launched from an instance store backed AMI is rebooted all the ephemeral storage data is still preserved.

NEW QUESTION 259

- (Topic 3)

A user has configured Auto Scaling with the minimum capacity as 2 and the desired capacity as 2. The user is trying to terminate one of the existing instance with the command:

```
as-terminate-instance-in-auto-scaling-group<Instance ID> --decrement-desired-capacity
```

What will Auto Scaling do in this scenario?

- A. Terminates the instance and does not launch a new instance
- B. Terminates the instance and updates the desired capacity to 1
- C. Terminates the instance and updates the desired capacity and minimum size to 1
- D. Throws an error

Answer: D

Explanation:

The Auto Scaling command `as-terminate-instance-in-auto-scaling-group <Instance ID>` will terminate the specific instance ID. The user is required to specify the parameter `--decrement-desired-capacity`. Then Auto Scaling will terminate the instance and decrease the desired capacity by 1. In this case since the minimum size is 2, Auto Scaling will not allow the desired capacity to go below 2. Thus, it will throw an error.

NEW QUESTION 262

- (Topic 3)

A user has enabled versioning on an S3 bucket. The user is using server side encryption for data at rest. If the user is supplying his own keys for encryption (SSE-C), what is recommended to the user for the purpose of security?

- A. The user should not use his own security key as it is not secure
- B. Configure S3 to rotate the user's encryption key at regular intervals
- C. Configure S3 to store the user's keys securely with SSL
- D. Keep rotating the encryption key manually at the client side

Answer: D

Explanation:

AWS S3 supports client side or server side encryption to encrypt all data at Rest. The server side encryption can either have the S3 supplied AES-256 encryption key or the user can send the key along with each API call to supply his own encryption key (SSE-C.. Since S3 does not store the encryption keys in SSE-C, it is recommended that the user should manage keys securely and keep rotating them regularly at the client side version.

NEW QUESTION 263

- (Topic 3)

A user is trying to create an EBS volume with the highest PIOPS supported by EBS. What is the minimum size of EBS required to have the maximum IOPS?

- A. 124
- B. 150
- C. 134
- D. 128

Answer: C

Explanation:

A provisioned IOPS EBS volume can range in size from 10 GB to 1 TB and the user can provision up to 4000 IOPS per volume. The ratio of IOPS provisioned to the volume size requested should be a maximum of 30.

NEW QUESTION 264

- (Topic 3)

A user has setup a custom application which generates a number in decimals. The user wants to track that number and setup the alarm whenever the number is above a certain limit. The application is sending the data to CloudWatch at regular intervals for this purpose. Which of the below mentioned statements is not true with respect to the above scenario?

- A. The user can get the aggregate data of the numbers generated over a minute and send it to CloudWatch
- B. The user has to supply the timezone with each data point
- C. CloudWatch will not truncate the number until it has an exponent larger than 126 (i.
- D. (1×10^{126}) .
- E. The user can create a file in the JSON format with the metric name and value and supply it to CloudWatch

Answer: B

NEW QUESTION 267

- (Topic 3)

A user has moved an object to Glacier using the life cycle rules. The user requests to restore the archive after 6 months. When the restore request is completed the user accesses that archive. Which of the below mentioned statements is not true in this condition?

- A. The archive will be available as an object for the duration specified by the user during the restoration request
- B. The restored object's storage class will be RRS
- C. The user can modify the restoration period only by issuing a new restore request with the updated period
- D. The user needs to pay storage for both RRS (restore
- E. and Glacier (Archiv
- F. Rates

Answer: B

Explanation:

AWS Glacier is an archival service offered by AWS. AWS S3 provides lifecycle rules to archive and restore objects from S3 to Glacier. Once the object is archived their storage class will change to Glacier. If the user sends a request for restore, the storage class will still be Glacier for the restored object. The user will be paying for both the archived copy as well as for the restored object. The object is available only for the duration specified in the restore request and if the user wants to modify that period, he has to raise another restore request with the updated duration.

NEW QUESTION 272

- (Topic 3)

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.

NEW QUESTION 273

- (Topic 3)

A user has configured an HTTPS listener on an ELB. The user has not configured any security policy which can help to negotiate SSL between the client and ELB. What will ELB do in this scenario?

- A. By default ELB will select the first version of the security policy
- B. By default ELB will select the latest version of the policy
- C. ELB creation will fail without a security policy
- D. It is not required to have a security policy since SSL is already installed

Answer: B

Explanation:

Elastic Load Balancing uses a Secure Socket Layer (SSL) negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. If the user has created an HTTPS/SSL listener without associating any security policy, Elastic Load Balancing will, by default, associate the latest version of the ELBSecurityPolicy-YYYY-MM with the load balancer.

NEW QUESTION 274

- (Topic 3)

A user has created a VPC with public and private subnets using the VPC wizard. Which of the below mentioned statements is not true in this scenario?

- A. The VPC will create a routing instance and attach it with a public subnet
- B. The VPC will create two subnets
- C. The VPC will create one internet gateway and attach it to VPC
- D. The VPC will launch one NAT instance with an elastic IP

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, the instances in the public subnet can receive inbound traffic directly from the internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance with an elastic IP. Wizard will also create two subnets with route tables. It will also create an internet gateway and attach it to the VPC.

NEW QUESTION 279

- (Topic 3)

A user runs the command "dd if=/dev/xvdf of=/dev/null bs=1M" on an EBS volume created from a snapshot and attached to a Linux instance. Which of the below mentioned activities is the user performing with the step given above?

- A. Pre warming the EBS volume
- B. Initiating the device to mount on the EBS volume
- C. Formatting the volume
- D. Copying the data from a snapshot to the device

Answer: A

Explanation:

When the user creates an EBS volume and is trying to access it for the first time it will encounter reduced IOPS due to wiping or initiating of the block storage. To avoid this as well as achieve the best performance it is required to pre warm the EBS volume. For a volume created from a snapshot and attached with a Linux OS, the "dd" command pre warms the existing data on EBS and any restored snapshots of volumes that have been previously fully pre warmed. This command maintains incremental snapshots; however, because this operation is read-only, it does not pre warm unused space that has never been written to on the original volume. In the command "dd if=/dev/xvdf of=/dev/null bs=1M", the parameter "if=input file" should be set to the drive that the user wishes to warm. The "of=output file" parameter should be set to the Linux null virtual device, /dev/null. The "bs" parameter sets the block size of the read operation; for optimal performance, this should be set to 1 MB.

NEW QUESTION 284

- (Topic 3)

A user is planning to schedule a backup for an EBS volume. The user wants security of the snapshot data. How can the user achieve data encryption with a snapshot?

- A. Use encrypted EBS volumes so that the snapshot will be encrypted by AWS
- B. While creating a snapshot select the snapshot with encryption
- C. By default the snapshot is encrypted by AWS
- D. Enable server side encryption for the snapshot using S3

Answer: A

Explanation:

AWS EBS supports encryption of the volume. It also supports creating volumes from existing snapshots provided the snapshots are created from encrypted volumes. The data at rest, the I/O as well as all the snapshots of the encrypted EBS will also be encrypted. EBS encryption is based on the AES-256 cryptographic algorithm, which is the industry standard.

NEW QUESTION 289

- (Topic 3)

A user has setup an EBS backed instance and attached 2 EBS volumes to it. The user has setup a CloudWatch alarm on each volume for the disk data. The user has stopped the EC2 instance and detached the EBS volumes. What will be the status of the alarms on the EBS volume?

- A. OK
- B. Insufficient Data
- C. Alarm
- D. The EBS cannot be detached until all the alarms are removed

Answer: B

Explanation:

Amazon CloudWatch alarm watches a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. Alarms invoke actions only for sustained state changes. There are three states of the alarm: OK, Alarm and Insufficient data. In this case since the EBS is detached and inactive the state will be Insufficient.

NEW QUESTION 290

- (Topic 3)

A sys admin is trying to understand the sticky session algorithm. Please select the correct sequence of steps, both when the cookie is present and when it is not, to help the admin understand the implementation of the sticky session:

ELB inserts the cookie in the response ELB chooses the instance based on the load balancing algorithm Check the cookie in the service request The cookie is found in the request The cookie is not found in the request

- A. 3,1,4,2 [Cookie is not Present] & 3,1,5,2 [Cookie is Present]
- B. 3,4,1,2 [Cookie is not Present] & 3,5,1,2 [Cookie is Present]
- C. 3,5,2,1 [Cookie is not Present] & 3,4,2,1 [Cookie is Present]
- D. 3,2,5,4 [Cookie is not Present] & 3,2,4,5 [Cookie is Present]

Answer: C

Explanation:

Generally AWS ELB routes each request to a zone with the minimum load. The Elastic Load Balancer provides a feature called sticky session which binds the user's session with a specific EC2 instance. The load balancer uses a special load-balancer-generated cookie to track the application instance for each request. When the load balancer receives a request, it first checks to see if this cookie is present in the request. If so, the request is sent to the application instance specified in the cookie. If there is no cookie, the load balancer chooses an application instance based on the existing load balancing algorithm. A cookie is inserted into the response for binding subsequent requests from the same user to that application instance.

NEW QUESTION 291

- (Topic 3)

A user has launched an EC2 instance. However, due to some reason the instance was terminated. If the user wants to find out the reason for termination, where can he find the details?

- A. It is not possible to find the details after the instance is terminated
- B. The user can get information from the AWS console, by checking the Instance description under the State transition reason label
- C. The user can get information from the AWS console, by checking the Instance description under the Instance Status Change reason label
- D. The user can get information from the AWS console, by checking the Instance description under the Instance Termination reason label

Answer: D

Explanation:

An EC2 instance, once terminated, may be available in the AWS console for a while after termination. The user can find the details about the termination from the description tab under the label State transition reason. If the instance is still running, there will be no reason listed. If the user has explicitly stopped or terminated the instance, the reason will be "User initiated shutdown".

NEW QUESTION 293

- (Topic 3)

A user is running a batch process on EBS backed EC2 instances. The batch process starts a few instances to process hadoop Map reduce jobs which can run between 50 – 600 minutes or sometimes for more time. The user wants to configure that the instance gets terminated only when the process is completed. How can the user configure this with CloudWatch?

- A. Setup the CloudWatch action to terminate the instance when the CPU utilization is less than 5%
- B. Setup the CloudWatch with Auto Scaling to terminate all the instances
- C. Setup a job which terminates all instances after 600 minutes
- D. It is not possible to terminate instances automatically

Answer: D

Explanation:

Amazon CloudWatch alarm watches a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup an action which terminates the instances when their CPU utilization is below a certain threshold for a certain period of time. The EC2 action can either terminate or stop the instance as part of the EC2 action.

NEW QUESTION 295

- (Topic 3)

A user has launched multiple EC2 instances for the purpose of development and testing in the same region. The user wants to find the separate cost for the production and development instances. How can the user find the cost distribution?

- A. The user should download the activity report of the EC2 services as it has the instance ID wise data
- B. It is not possible to get the AWS cost usage data of single region instances separately
- C. The user should use Cost Distribution Metadata and AWS detailed billing
- D. The user should use Cost Allocation Tags and AWS billing reports

Answer: D

Explanation:

AWS provides cost allocation tags to categorize and track the AWS costs. When the user applies tags to his AWS resources (such as Amazon EC2 instances or Amazon S3 buckets), AWS generates a cost allocation report as a comma-separated value (CSV file) with the usage and costs aggregated by those tags. The user can apply tags which represent business categories (such as cost centres, application names, or instance type – Production/Dev. to organize usage costs across multiple services.

NEW QUESTION 300

- (Topic 3)

A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered. What time out period should the user specify for connection draining?

- A. 5 minutes
- B. 1 hour
- C. 30 minutes
- D. 2 hours

Answer: B

NEW QUESTION 303

- (Topic 3)

A user has setup an Auto Scaling group. The group has failed to launch a single instance for more than 24 hours. What will happen to Auto Scaling in this condition?

- A. Auto Scaling will keep trying to launch the instance for 72 hours
- B. Auto Scaling will suspend the scaling process
- C. Auto Scaling will start an instance in a separate region
- D. The Auto Scaling group will be terminated automatically

Answer: B

Explanation:

If Auto Scaling is trying to launch an instance and if the launching of the instance fails continuously, it will suspend the processes for the Auto Scaling groups since it repeatedly failed to launch an instance. This is known as an administrative suspension. It commonly applies to the Auto Scaling group that has no running instances which is trying to launch instances for more than 24 hours, and has not succeeded in that to do so.

NEW QUESTION 305

- (Topic 3)

In order to optimize performance for a compute cluster that requires low inter-node latency, which feature in the following list should you use?

- A. AWS Direct Connect
- B. Placement Groups
- C. VPC private subnets
- D. EC2 Dedicated Instances
- E. Multiple Availability Zones

Answer: D

NEW QUESTION 307

- (Topic 3)

A storage admin wants to encrypt all the objects stored in S3 using server side encryption. The user does not want to use the AES 256 encryption key provided by S3. How can the user achieve this?

- A. The admin should upload his secret key to the AWS console and let S3 decrypt the objects
- B. The admin should use CLI or API to upload the encryption key to the S3 bucket
- C. When making a call to the S3 API mention the encryption key URL in each request
- D. S3 does not support client supplied encryption keys for server side encryption
- E. The admin should send the keys and encryption algorithm with each API call

Answer: D

Explanation:

AWS S3 supports client side or server side encryption to encrypt all data at rest. The server side encryption can either have the S3 supplied AES-256 encryption key or the user can send the key along with each API call to supply his own encryption key. Amazon S3 never stores the user's encryption key. The user has to supply it for each encryption or decryption call.

NEW QUESTION 308

- (Topic 3)

A .NET application that you manage is running in Elastic Beanstalk. Your developers tell you they will need access to application log files to debug issues that arise. The infrastructure will scale up and down.

How can you ensure the developers will be able to access only the log files?

- A. Access the log files directly from Elastic Beanstalk
- B. Enable log file rotation to S3 within the Elastic Beanstalk configuration
- C. Ask your developers to enable log file rotation in the applications web.config file
- D. Connect to each Instance launched by Elastic Beanstalk and create a Windows Scheduled task to rotate the log files to S3.

Answer: D

Explanation:

Reference:

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.loggingS3.title.html>

NEW QUESTION 310

- (Topic 3)

A user has enabled versioning on an S3 bucket. The user is using server side encryption for data at Rest. If the user is supplying his own keys for encryption (SSE-C), which of the below mentioned statements is true?

- A. The user should use the same encryption key for all versions of the same object
- B. It is possible to have different encryption keys for different versions of the same object
- C. AWS S3 does not allow the user to upload his own keys for server side encryption
- D. The SSE-C does not work when versioning is enabled

Answer: B

Explanation:

AWS S3 supports client side or server side encryption to encrypt all data at rest. The server side encryption can either have the S3 supplied AES-256 encryption key or the user can send the key along with each API call to supply his own encryption key (SSE-C). If the bucket is versioning-enabled, each object version uploaded by the user using the SSE-C feature can have its own encryption key. The user is responsible for tracking which encryption key was used for which object's version

NEW QUESTION 312

- (Topic 3)

A sysadmin has created the below mentioned policy on an S3 bucket named cloudacademy. What does this policy define?

```
"Statement": [{
  "Sid": "Stmnt1388811069831",
  "Effect": "Allow",
  "Principal": { "AWS": "*" },
  "Action": [ "s3:GetObjectAcl", "s3:ListBucket" ],
  "Resource": [ "arn:aws:s3:::cloudacademy" ]
}]
```

- A. It will make the cloudacademy bucket as well as all its objects as public
- B. It will allow everyone to view the ACL of the bucket
- C. It will give an error as no object is defined as part of the policy while the action defines the rule about the object
- D. It will make the cloudacademy bucket as public

Answer: D

Explanation:

A sysadmin can grant permission to the S3 objects or the buckets to any user or make objects public using the bucket policy and user policy. Both use the JSON-based access policy language. Generally if the user is defining the ACL on the bucket, the objects in the bucket do not inherit it and vice versa. The bucket policy can be defined at the bucket level which allows the objects as well as the bucket to be public with a single policy applied to that bucket. In the sample policy the action says "S3:ListBucket" for effect Allow on Resource arn:aws:s3:::cloudacademy. This will make the cloudacademy bucket public.

```
"Statement": [{
  "Sid": "Stmnt1388811069831",
  "Effect": "Allow",
  "Principal": { "AWS": "*" },
  "Action": [ "s3:GetObjectAcl", "s3:ListBucket" ],
  "Resource": [ "arn:aws:s3:::cloudacademy" ]
}]
```

NEW QUESTION 317

- (Topic 3)

A user has enabled termination protection on an EC2 instance. The user has also set Instance initiated shutdown behaviour to terminate. When the user shuts down the instance from the OS, what will happen?

- A. The OS will shutdown but the instance will not be terminated due to protection
- B. It will terminate the instance
- C. It will not allow the user to shutdown the instance from the OS
- D. It is not possible to set the termination protection when an Instance initiated shutdown is set to Terminate

Answer: B

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 user can also setup shutdown behaviour for an EBS backed instance to guide the instance on what should be done when he initiates shutdown from the OS using Instance initiated shutdown behaviour. If the instance initiated behaviour is set to terminate and the user shuts off the OS even though termination protection is enabled, it will still terminate the instance.

NEW QUESTION 322

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

AWS-SysOps Practice Exam Features:

- * AWS-SysOps Questions and Answers Updated Frequently
- * AWS-SysOps Practice Questions Verified by Expert Senior Certified Staff
- * AWS-SysOps Most Realistic Questions that Guarantee you a Pass on Your First Try
- * AWS-SysOps Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-SysOps Practice Test Here](#)