

Amazon

Exam Questions AWS-SysOps

AWS Certified SysOps Administrator Associate



NEW QUESTION 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 CloudWatch alarm that notifies your operations team should the primary monitoring instance become unhealthy
- 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 instance
- C. Set a CloudWatch alarm based on the CPU utilization of the monitoring instance and have the alarm notify your operations team if the CPU usage exceeds 50% for 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 SQS 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 another backup monitoring instance and assume (the role of the previous monitoring instance and begin adding messages to the SQS queue

Answer: D

NEW QUESTION 2

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 3

When preparing for a compliance assessment of your system built inside of AWS. what are three best-practices for you to prepare for an audit? Choose 3 answers

- A. Gather evidence of your IT operational controls
- B. Request and obtain applicable third-party audited AWS compliance reports and certifications
- C. Request and obtain a compliance and security tour of an AWS data center for a pre-assessment security review
- D. Request and obtain approval from AWS to perform relevant network scans and in-depth penetration tests of your system's Instances and endpoints
- E. Schedule meetings with AWS's third-party auditors to provide evidence of AWS compliance that maps to your control objectives

Answer: ABD

NEW QUESTION 4

You are managing a legacy application inside VPC with hard coded IP addresses in its configuration.

Which two mechanisms will allow the application to failover to new instances without the need for reconfiguration? Choose 2 answers

- A. Create an ELB to reroute traffic to a failover instance
- B. Create a secondary ENI that can be moved to a failover instance
- C. Use Route53 health checks to fail traffic over to a failover instance
- D. Assign a secondary private IP address to the primary ENI that can be moved to a failover instance

Answer: AD

NEW QUESTION 5

You have decided to change the Instance type for instances running in your application tier that are using Auto Scaling. In which area below would you change the instance type definition?

- A. Auto Scaling launch configuration
- B. Auto Scaling group
- C. Auto Scaling policy
- D. Auto Scaling tags

Answer: A

Explanation: Reference:

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/WhatIsAutoScaling.html>

NEW QUESTION 6

You are designing a system that has a Bastion host. This component needs to be highly available without human intervention. Which of the following approaches would you select?

- A. Run the bastion on two instances one in each AZ

- B. Run the bastion on an active Instance in one AZ and have an AMI ready to boot up in the event of failure
- C. Configure the bastion instance in an Auto Scaling group Specify the Auto Scaling group to include multiple AZs but have a min-size of 1 and max-size of 1
- D. Configure an ELB in front of the bastion instance

Answer: C

NEW QUESTION 7

You have set up Individual AWS accounts for each project. You have been asked to make sure your AWS Infrastructure costs do not exceed the budget set per project for each month.

Which of the following approaches can help ensure that you do not exceed the budget each month?

- A. Consolidate your accounts so you have a single bill for all accounts and projects
- B. Set up auto scaling with CloudWatch alarms using SNS to notify you when you are running too many Instances in a given account
- C. Set up CloudWatch billing alerts for all AWS resources used by each project, with a notification occurring when the amount for each resource tagged to a particular project matches the budget allocated to the projec
- D. Set up CloudWatch billing alerts for all AWS resources used by each account, with email notifications when it hits 50%. 80% and 90% of its budgeted monthly spend

Answer: C

NEW QUESTION 8

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

You are running a web-application on AWS consisting of the following components an Elastic Load Balancer (ELB) an Auto-Scaling Group of EC2 instances running Linux/PHP/Apache, and Relational DataBase Service (RDS) MySQL.

Which security measures fall into AWS's responsibility?

- A. Protect the EC2 instances against unsolicited access by enforcing the principle of least-privilege access
- B. Protect against IP spoofing or packet sniffing
- C. Assure all communication between EC2 instances and ELB is encrypted
- D. Install latest security patches on EL
- E. RDS and EC2 instances

Answer: B

NEW QUESTION 10

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

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 15

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 cause 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 data
- D. Use SQS to queue up the requests for the technical posts and deliver them out of the queue
- E. Use Route53 health checks to fail over to an S3 bucket for an error page

Answer: ACE

NEW QUESTION 17

What is a placement group?

- A. A collection of Auto Scaling groups in the same Region
- B. Feature that enables EC2 instances to interact with each other via high bandwidth, low latency connections
- C. A collection of Elastic Load Balancers in the same Region or Availability Zone
- D. A collection of authorized CloudFront edge locations for a distribution

Answer: B

Explanation: Reference:

<http://aws.amazon.com/ec2/faqs/>

NEW QUESTION 19

An application that you are managing has EC2 instances & DynamoDB tables deployed to several AWS Regions. In order to monitor the performance of the application globally, you would like to see two graphs: 1) Avg CPU Utilization across all EC2 instances and 2) Number of Throttled Requests for all DynamoDB tables.

How can you accomplish this?

- A. Tag your resources with the application name, and select the tag name as the dimension in the CloudWatch Management console to view the respective graphs
- B. Use the CloudWatch CLI tools to pull the respective metrics from each regional endpoint. Aggregate the data offline & store it for graphing in CloudWatch
- C. Add SNMP traps to each instance and DynamoDB table. Leverage a central monitoring server to capture data from each instance and table. Put the aggregate data into CloudWatch for graphing
- D. Add a CloudWatch agent to each instance and attach one to each DynamoDB table
- E. When configuring the agent, set the appropriate application name & view the graphs in CloudWatch

Answer: C

NEW QUESTION 23

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

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

Answer: C

Explanation: Reference:

<http://aws.amazon.com/cloudwatch/>

NEW QUESTION 25

Which of the following are characteristics of Amazon VPC subnets?

Choose 2 answers

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

Answer: CE

NEW QUESTION 27

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

Which two things do we need to confirm in the VPC settings so that these EC2 instances can communicate inside the VPC? Choose 2 answers

- A. A network ACL that allows communication between the two subnets

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

Answer: AD

NEW QUESTION 32

Your company is moving towards tracking web page users with a small tracking image loaded on each page. Currently you are serving this image out of US-East, but are starting to get concerned about the time it takes to load the image for users on the west coast.

What are the two best ways to speed up serving this image?

Choose 2 answers

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

Answer: AD

NEW QUESTION 35

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 37

Which two AWS services provide out-of-the-box user configurable automatic backup-as-a-service and backup rotation options?

Choose 2 answers

- A. Amazon S3
- B. Amazon RDS
- C. Amazon EBS
- D. Amazon Redshift

Answer: BD

NEW QUESTION 38

You are using ElastiCache Memcached to store session state and cache database queries in your infrastructure. You notice in CloudWatch that Evictions and GetMisses are both very high.

What two actions could you take to rectify this?

Choose 2 answers

- A. Increase the number of nodes in your cluster
- B. Tweak the max_item_size parameter
- C. Shrink the number of nodes in your cluster
- D. Increase the size of the nodes in the cluster

Answer: BD

NEW QUESTION 41

An organization's security policy requires multiple copies of all critical data to be replicated across at least a primary and backup data center. The organization has decided to store some critical data on Amazon S3.

Which option should you implement to ensure this requirement is met?

- A. Use the S3 copy API to replicate data between two S3 buckets in different regions
- B. You do not need to implement anything since S3 data is automatically replicated between regions
- C. Use the S3 copy API to replicate data between two S3 buckets in different facilities within an AWS Region
- D. You do not need to implement anything since S3 data is automatically replicated between multiple facilities within an AWS Region

Answer: D

NEW QUESTION 44

Which services allow the customer to retain full administrative privileges of the underlying EC2 instances?

Choose 2 answers

- A. Amazon Elastic MapReduce

- B. Elastic Load Balancing
- C. AWS Elastic Beanstalk
- D. Amazon ElastiCache
- E. Amazon Relational Database service

Answer: AC

NEW QUESTION 49

You have been asked to propose a multi-region deployment of a web-facing application where a controlled portion of your traffic is being processed by an alternate region.

Which configuration would achieve that goal?

- A. Route53 record sets with weighted routing policy
- B. Route53 record sets with latency based routing policy
- C. Auto Scaling with scheduled scaling actions set
- D. Elastic Load Balancing with health checks enabled

Answer: D

Explanation: Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html>

NEW QUESTION 54

When attached to an Amazon VPC which two components provide connectivity with external networks? Choose 2 answers

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

Answer: CD

NEW QUESTION 57

An organization has configured a VPC with an Internet Gateway (IGW), pairs of public and private subnets (each with one subnet per Availability Zone), and an Elastic Load Balancer (ELB) configured to use the public subnets. The application's web tier leverages the ELB. Auto Scaling and a multi-AZ RDS database instance. The organization would like to eliminate any potential single points of failure in this design. What step should you take to achieve this organization's objective?

- A. Nothing, there are no single points of failure in this architecture
- B. Create and attach a second IGW to provide redundant internet connectivity
- C. Create and configure a second Elastic Load Balancer to provide a redundant load balance
- D. Create a second multi-AZ RDS instance in another Availability Zone and configure replication to provide a redundant database

Answer: A

NEW QUESTION 58

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 62

When creation of an EBS snapshot is initiated but not completed the EBS volume?

- A. Cannot be detached or attached to an EC2 instance until the snapshot completes
- B. Can be used in read-only mode while the snapshot is in progress
- C. Can be used while the snapshot is in progress
- D. Cannot be used until the snapshot completes

Answer: C

Explanation: Reference:

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

NEW QUESTION 65

A user is planning to setup infrastructure on AWS for the Christmas sales. The user is planning to use Auto Scaling based on the schedule for proactive scaling. What advice would you give to the user?

- A. It is good to schedule now because if the user forgets later on it will not scale up
- B. The scaling should be setup only one week before Christmas
- C. Wait till end of November before scheduling the activity
- D. It is not advisable to use scheduled based scaling

Answer: C

Explanation:

Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. The user can specify any date in the future to scale up or down during that period. As per Auto Scaling the user can schedule an action for up to a month in the future. Thus, it is recommended to wait until end of November before scheduling for Christmas.

NEW QUESTION 66

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 69

A user wants to disable connection draining on an existing ELB. Which of the below mentioned statements helps the user disable connection draining on the ELB?

- A. The user can only disable connection draining from CLI
- B. It is not possible to disable the connection draining feature once enabled
- C. The user can disable the connection draining feature from EC2 -> ELB console or from CLI
- D. The user needs to stop all instances before disabling connection draining

Answer: C

Explanation:

The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that inflight requests continue to be served. The user can enable or disable connection draining from the AWS EC2 console -> ELB or using CLI.

NEW QUESTION 74

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 role
- 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 79

A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created a public subnet CIDR (20.0.0.0/24) and VPN only subnets CIDR (20.0.1.0/24) along with the VPN gateway (vgw-12345) to connect to the user's data centre. Which of the below mentioned options is a valid entry for the main route table in this scenario?

- A. Destination: 20.0.0.0/24 and Target: vgw-12345
- B. Destination: 20.0.0.0/16 and Target: ALL
- C. Destination: 20.0.1.0/16 and Target: vgw-12345
- D. Destination: 0.0.0.0/0 and Target: vgw-12345

Answer: D

Explanation:

The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all traffic of the VPN subnet. Here are the valid entries for the main route table in this scenario: Destination: 0.0.0.0/0 & Target: vgw-12345 (To route all internet traffic to the VPN gateway. Destination: 20.0.0.0/16 & Target: local (To allow local routing in VPC.

NEW QUESTION 80

An organization is generating digital policy files which are required by the admins for verification. Once the files are verified they may not be required in the future unless there is some compliance issue. If the organization wants to save them in a cost effective way, which is the best possible solution?

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

Answer: D

Explanation:

Amazon S3 stores objects according to their storage class. There are three major storage classes: Standard, Reduced Redundancy and Glacier. Standard is for AWS S3 and provides very high durability. However, the costs are a little higher. Reduced redundancy is for less critical files. Glacier is for archival and the files which are accessed infrequently. It is an extremely low-cost storage service that provides secure and durable storage for data archiving and backup.

NEW QUESTION 82

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 83

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 86

A user has received a message from the support team that an issue occurred 1 week back between 3 AM to 4 AM and the EC2 server was not reachable. The user is checking the CloudWatch metrics of that instance. How can the user find the data easily using the CloudWatch console?

- A. The user can find the data by giving the exact values in the time Tab under CloudWatch metrics
- B. The user can find the data by filtering values of the last 1 week for a 1 hour period in the Relative tab under CloudWatch metrics
- C. It is not possible to find the exact time from the consol
- D. The user has to use CLI to provide the specific time
- E. The user can find the data by giving the exact values in the Absolute tab under CloudWatch metrics

Answer: D

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.

NEW QUESTION 90

A sys admin is maintaining an application on AWS. The application is installed on EC2 and user has configured ELB and Auto Scaling. Considering future load

increase, the user is planning to launch new servers proactively so that they get registered with ELB. How can the user add these instances with Auto Scaling?

- A. Increase the desired capacity of the Auto Scaling group
- B. Increase the maximum limit of the Auto Scaling group
- C. Launch an instance manually and register it with ELB on the fly
- D. Decrease the minimum limit of the Auto Scaling group

Answer: A

Explanation:

A user can increase the desired capacity of the Auto Scaling group and Auto Scaling will launch a new instance as per the new capacity. The newly launched instances will be registered with ELB if Auto Scaling group is configured with ELB. If the user decreases the minimum size the instances will be removed from Auto Scaling. Increasing the maximum size will not add instances but only set the maximum instance cap.

NEW QUESTION 93

A sys admin has created the below mentioned policy and applied to an S3 object named aws.jpg. The aws.jpg is inside a bucket named cloudacademy. What does this policy define?

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

- A. It is not possible to define a policy at the object level
- B. It will make all the objects of the bucket cloudacademy as public
- C. It will make the bucket cloudacademy as public
- D. the aws.jpg object as public

Answer: A

Explanation:

A system admin can grant permission to the S3 objects or 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. It cannot be applied at the object level.

NEW QUESTION 98

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 100

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 101

A user is planning to use AWS Cloud formation for his automatic deployment requirements. Which of the below mentioned components are required as a part of the template?

- A. Parameters
- B. Outputs
- C. Template version
- D. Resources

Answer: D

Explanation:

AWS Cloud formation is an application management tool which provides application modelling, deployment, configuration, management and related activities. The template is a JSON-format, text-based file that describes all the AWS resources required to deploy and run an application. It can have option fields, such as Template Parameters, Output, Data tables, and Template file format version. The only mandatory value is Resource. The user can define the AWS services which will be used/ created by this template inside the Resource section

NEW QUESTION 104

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 108

A user has created a queue named "myqueue" in US-East region with AWS SQS. The user's AWS account ID is 123456789012. If the user wants to perform some action on this queue, which of the below Queue URL should he use?

- A. <http://sqs.us-east-1.amazonaws.com/123456789012/myqueue>
- B. <http://sqs.amazonaws.com/123456789012/myqueue>
- C. <http://sq>
- D. 123456789012.us-east-1.amazonaws.com/myqueue
- E. [http:// 123456789012.sq](http://123456789012.sq)
- F. us-east-1.amazonaws.com/myqueue

Answer: A

Explanation:

When creating a new queue in SQS, the user must provide a queue name that is unique within the scope of all queues of user's account. If the user creates queues using both the latest WSDL and a previous version, he will have a single namespace for all his queues. Amazon SQS assigns each queue created by user an identifier called a queue URL, which includes the queue name and other components that Amazon SQS determines. Whenever the user wants to perform an action on a queue, he must provide its queue URL. The queue URL for the account id 123456789012 & queue name "myqueue" in US-East-1 region will be <http://sqs.us-east-1.amazonaws.com/123456789012/myqueue>.

NEW QUESTION 111

An organization is planning to use AWS for 5 different departments. The finance department is responsible to pay for all the accounts. However, they want the cost separation for each account to map with the right cost centre. How can the finance department achieve this?

- A. Create 5 separate accounts and make them a part of one consolidate billing
- B. Create 5 separate accounts and use the IAM cross account access with the roles for better management
- C. Create 5 separate IAM users and set a different policy for their access
- D. Create 5 separate IAM groups and add users as per the department's employees

Answer: A

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.

NEW QUESTION 113

A user has launched an ELB which has 5 instances registered with it. The user deletes the ELB by mistake. What will happen to the instances?

- A. ELB will ask the user whether to delete the instances or not
- B. Instances will be terminated
- C. ELB cannot be deleted if it has running instances registered with it
- D. Instances will keep running

Answer: D

Explanation:

When the user deletes the Elastic Load Balancer, all the registered instances will be deregistered. However, they will continue to run. The user will incur charges if he does not take any action on those instances.

NEW QUESTION 118

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

- A. It will not allow the user to create the private subnet due to a CIDR overlap
- B. It will allow the user to create a private subnet with CIDR as 20.0.0.128/25
- C. This statement is wrong as AWS does not allow CIDR 20.0.0.0/25
- D. It will not allow the user to create a private subnet due to a wrong CIDR range

Answer: B

Explanation:

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

NEW QUESTION 119

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 120

An organization has setup consolidated billing with 3 different AWS accounts. Which of the below mentioned advantages will organization receive in terms of the AWS pricing?

- A. The consolidated billing does not bring any cost advantage for the organization
- B. All AWS accounts will be charged for S3 storage by combining the total storage of each account
- C. The EC2 instances of each account will receive a total of 750*3 micro instance hours free
- D. The free usage tier for all the 3 accounts will be 3 years and not a single year

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. For billing purposes, AWS treats all the accounts on the consolidated bill as one account. Some services, such as Amazon EC2 and Amazon S3 have volume pricing tiers across certain usage dimensions that give the user lower prices when he uses the service more.

NEW QUESTION 121

A user has setup connection draining with ELB to allow in-flight requests to continue while the instance is being deregistered through Auto Scaling. If the user has not specified the draining time, how long will ELB allow inflight requests traffic to continue?

- A. 600 seconds
- B. 3600 seconds
- C. 300 seconds
- D. 0 seconds

Answer: C

Explanation:

The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that inflight requests continue to be served. The user can specify a maximum time (3600 seconds. for the load balancer to keep the connections alive before reporting the instance as deregistered. If the user does not specify the maximum timeout period, by default, the load balancer will close the connections to the deregistering instance after 300 seconds.

NEW QUESTION 124

A user has created an S3 bucket which is not publicly accessible. The bucket is having thirty objects which are also private. If the user wants to make the objects public, how can he configure this with minimal efforts?

- A. The user should select all objects from the console and apply a single policy to mark them public
- B. The user can write a program which programmatically makes all objects public using S3 SDK
- C. Set the AWS bucket policy which marks all objects as public
- D. Make the bucket ACL as public so it will also mark all objects as public

Answer: C

Explanation:

A system admin can grant permission of the S3 objects or buckets to any user or make the 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.

NEW QUESTION 128

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

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

Answer: B

Explanation:

SQS allows the user to move data between distributed components of applications so they can perform different tasks without losing messages or requiring each component to be always available. The user can delete a queue at any time, whether it is empty or not. It is important to note that queues retain messages for a set period of time. By default, a queue retains messages for four days.

NEW QUESTION 132

A user has created a web application with Auto Scaling. The user is regularly monitoring the application and he observed that the traffic is highest on Thursday and Friday between 8 AM to 6 PM. What is the best solution to handle scaling in this case?

- A. Add a new instance manually by 8 AM Thursday and terminate the same by 6 PM Friday
- B. Schedule Auto Scaling to scale up by 8 AM Thursday and scale down after 6 PM on Friday
- C. Schedule a policy which may scale up every day at 8 AM and scales down by 6 PM
- D. Configure a batch process to add a instance by 8 AM and remove it by Friday 6 PM

Answer: B

Explanation:

Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. In this case the load increases by Thursday and decreases by Friday. Thus, the user can setup the scaling activity based on the predictable traffic patterns of the web application using Auto Scaling scale by Schedule.

NEW QUESTION 137

A user has setup a CloudWatch alarm on an EC2 action when the CPU utilization is above 75%. The alarm sends a notification to SNS on the alarm state. If the user wants to simulate the alarm action how can he achieve this?

- A. Run activities on the CPU such that its utilization reaches above 75%
- B. From the AWS console change the state to 'Alarm'
- C. The user can set the alarm state to 'Alarm' using CLI
- D. Run the SNS action manually

Answer: C

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 test an alarm by setting it to any state using the SetAlarmState API (mon-set-alarm-state command). This temporary state change lasts only until the next alarm comparison occurs.

NEW QUESTION 138

A user has enabled detailed CloudWatch metric monitoring on an Auto Scaling group. Which of the below mentioned metrics will help the user identify the total number of instances in an Auto Scaling group including pending, terminating and running instances?

- A. GroupTotalInstances
- B. GroupSumInstances
- C. It is not possible to get a count of all the three metrics together
- D. The user has to find the individual number of running, terminating and pending instances and sum it
- E. GroupInstancesCount

Answer: A

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. For Auto Scaling, CloudWatch provides various metrics to get the group information, such as the Number of Pending, Running or Terminating instances at any moment. If the user wants to get the total number of Running, Pending and Terminating instances at any moment, he can use the GroupTotalInstances metric.

NEW QUESTION 143

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 148

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 150

A user has setup a web application on EC2. The user is generating a log of the application performance at every second. There are multiple entries for each second. If the user wants to send that data to CloudWatch every minute, what should he do?

- A. The user should send only the data of the 60th second as CloudWatch will map the receive data timezone with the sent data timezone
- B. It is not possible to send the custom metric to CloudWatch every minute
- C. Give CloudWatch the Min, Max, Sum, and SampleCount of a number of every minute
- D. Calculate the average of one minute and send the data to CloudWatch

Answer: C

Explanation:

Amazon CloudWatch aggregates statistics according to the period length that the user has specified while getting data from CloudWatch. The user can publish as many data points as he wants with the same or similar time stamps. CloudWatch aggregates them by the period length when the user calls get statistics about those data points. CloudWatch records the average (sum of all items divided by the number of items. of the values received for every 1-minute period, as well as the number of samples, maximum value, and minimum value for the same time period. CloudWatch will aggregate all the data which have time stamps within a one-minute period.

NEW QUESTION 154

A user has launched an EBS backed instance. The user started the instance at 9 AM in the morning. Between 9 AM to 10 AM, the user is testing some script. Thus, he stopped the instance twice and restarted it. In the same hour the user rebooted the instance once. For how many instance hours will AWS charge the user?

- A. 3 hours
- B. 4 hours
- C. 2 hours
- D. 1 hour

Answer: A

Explanation:

A user can stop/start or 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. When the instance is rebooted AWS will not charge the user for the extra hours. In case the user stops the instance, AWS does not charge the running cost but charges only the EBS storage cost. If the user starts and stops the instance multiple times in a single hour, AWS will charge the user for every start and stop. In this case, since the instance was rebooted twice, it will cost the user for 3 instance hours.

NEW QUESTION 159

A user has configured the Auto Scaling group with the minimum capacity as 3 and the maximum capacity as 5. When the user configures the AS group, how many instances will Auto Scaling launch?

- A. 3
- B. 0
- C. 5
- D. 2

Answer: C

NEW QUESTION 160

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 162

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 163

An organization, which has the AWS account ID as 999988887777, has created 50 IAM users. All the users are added to the same group cloudacademy. If the organization has enabled that each IAM user can login with the AWS console, which AWS login URL will the IAM users use?

- A. [https:// 999988887777.signin.aws.amazon.com/console/](https://999988887777.signin.aws.amazon.com/console/)
- B. [https:// signin.aws.amazon.com/cloudacademy/](https://signin.aws.amazon.com/cloudacademy/)
- C. [https:// cloudacademy.signin.aws.amazon.com/999988887777/console/](https://cloudacademy.signin.aws.amazon.com/999988887777/console/)
- D. [https:// 999988887777.aws.amazon.com/ cloudacademy/](https://999988887777.aws.amazon.com/cloudacademy/)

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. Once the organization has created the IAM users, they will have a separate AWS console URL to login to the AWS console. The console login URL for the IAM user will be [https:// AWS_Account_ID.signin.aws.amazon.com/console/](https://AWS_Account_ID.signin.aws.amazon.com/console/). It uses only the AWS account ID and does not depend on the group or user ID.

NEW QUESTION 167

A user has configured an Auto Scaling group with ELB. The user has enabled detailed CloudWatch monitoring on Elastic Load balancing. Which of the below mentioned statements will help the user understand this functionality better?

- A. ELB sends data to CloudWatch every minute only and does not charge the user
- B. ELB will send data every minute and will charge the user extra
- C. ELB is not supported by CloudWatch
- D. It is not possible to setup detailed monitoring for ELB

Answer: A

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. Elastic Load Balancing includes 10 metrics and 2 dimensions, and sends data to CloudWatch every minute. This does not cost extra.

NEW QUESTION 169

A user has created a subnet with VPC and launched an EC2 instance in that subnet with only default settings. Which of the below mentioned options is ready to use on the EC2 instance as soon as it is launched?

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

Answer: B

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to a user's AWS account. A subnet is a range of IP addresses in the VPC. The user can launch the AWS resources into a subnet. There are two supported platforms into which a user can launch instances: EC2-Classic and EC2-VPC. When the user launches an instance which is not a part of the non-default subnet, it will only have a private IP assigned to it. The instances part of a subnet can communicate with each other but cannot communicate over the internet or to the AWS services, such as RDS / S3.

NEW QUESTION 174

A user is trying to delete an Auto Scaling group from CLI. Which of the below mentioned steps are to be performed by the user?

- A. Terminate the instances with the `ec2-terminate-instance` command
- B. Terminate the Auto Scaling instances with the `as-terminate-instance` command
- C. Set the minimum size and desired capacity to 0
- D. There is no need to change the capacity
- E. Run the `as-delete-group` command and it will reset all values to 0

Answer: C

Explanation:

If the user wants to delete the Auto Scaling group, the user should manually set the values of the minimum and desired capacity to 0. Otherwise Auto Scaling will not allow for the deletion of the group from CLI. While trying from the AWS console, the user need not set the values to 0 as the Auto Scaling console will automatically do so.

NEW QUESTION 176

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

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

Answer: C

Explanation:

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

NEW QUESTION 179

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 180

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 182

A user is publishing custom metrics to CloudWatch. Which of the below mentioned statements will help the user understand the functionality better?

- A. The user can use the CloudWatch Import tool
- B. The user should be able to see the data in the console after around 15 minutes
- C. If the user is uploading the custom data, the user must supply the namespace, timezone, and metric name as part of the command
- D. The user can view as well as upload data using the console, CLI and APIs

Answer: B

Explanation:

AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. The user has to always include the namespace as a part of the request. However, the other parameters are optional. If the user has uploaded data using CLI, he can view it as a graph inside the console. The data will take around 2 minutes to upload but can be viewed only after around 15 minutes.

NEW QUESTION 185

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

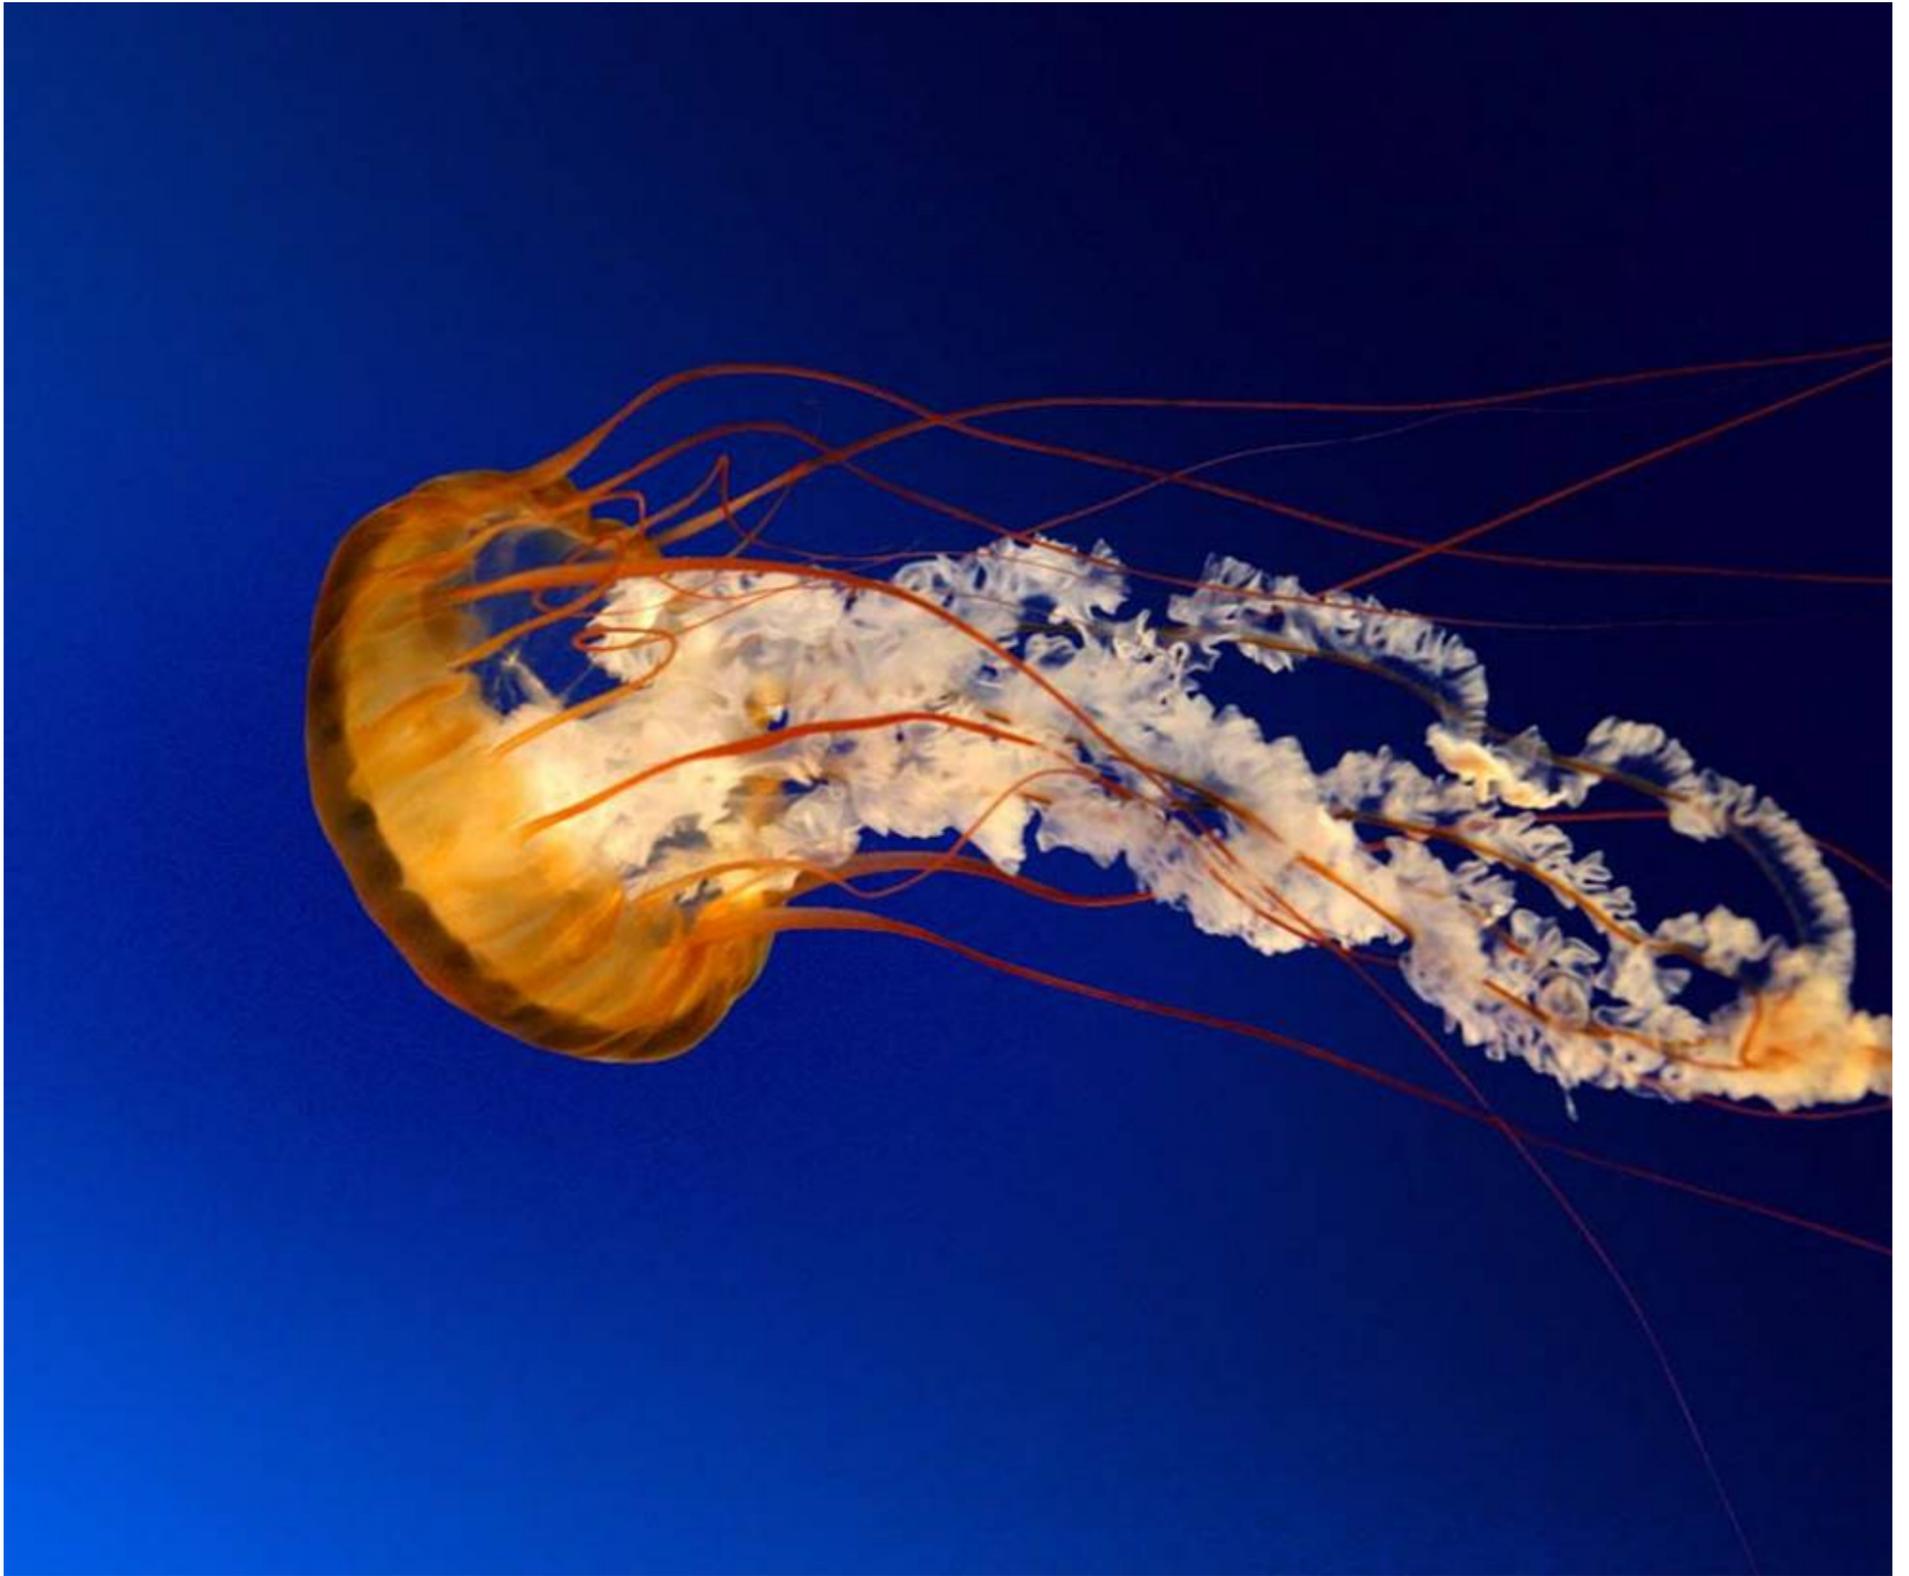
Explanation:

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 190

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 192

A user is trying to understand AWS SNS. To which of the below mentioned end points is SNS unable to send a notification?

- A. Email JSON
- B. HTTP
- C. AWS SQS
- D. AWS SES

Answer: D

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 select one of the following transports as part of the subscription requests: "HTTP", "HTTPS", "Email", "Email-JSON", "SQS", "and SMS".

NEW QUESTION 193

A user has configured a VPC with a new subnet. The user has created a security group. The user wants to configure that instances of the same subnet communicate with each other. How can the user configure this with the security group?

- A. There is no need for a security group modification as all the instances can communicate with each other inside the same subnet
- B. Configure the subnet as the source in the security group and allow traffic on all the protocols and ports
- C. Configure the security group itself as the source and allow traffic on all the protocols and ports
- D. The user has to use VPC peering to configure this

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. Security groups work at the instance level. If the user is using the default security group it will have a rule which allows the instances to communicate with other. For a new security group the user has to specify the rule, add it to define the source as the security group itself, and select all the protocols and ports for that source.

NEW QUESTION 198

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 AWS
- 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 202

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 Provisioned IOPS EBS
- B. Launch the EC2 instance as EBS enhanced with Provisioned IOPS EBS
- C. Launch the EC2 instance as EBS dedicated with Provisioned IOPS EBS
- D. Launch the EC2 instance as EBS optimized with Provisioned IOPS 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 203

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

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

Answer: D

Explanation: Reference:

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

NEW QUESTION 206

A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling terminate process only for a while. What will happen to the availability zone rebalancing process (AZRebalance) during this period?

- A. Auto Scaling will not launch or terminate any instances
- B. Auto Scaling will allow the instances to grow more than the maximum size
- C. Auto Scaling will keep launching instances till the maximum instance size
- D. It is not possible to suspend the terminate process while keeping the launch active

Answer: B

Explanation:

Auto Scaling performs various processes, such as Launch, Terminate, Availability Zone Rebalance (AZRebalance. etc. The AZRebalance process type seeks to maintain a balanced number of instances across Availability Zones within a region. If the user suspends the Terminate process, the AZRebalance process can cause the Auto Scaling group to grow up to ten percent larger than the maximum size. This is because Auto Scaling allows groups to temporarily grow larger than the maximum size during rebalancing activities. If Auto Scaling cannot terminate instances, the Auto Scaling group could remain up to ten percent larger than the maximum size until the user resumes the Terminate process type.

NEW QUESTION 208

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 211

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

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

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. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:

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

NEW QUESTION 216

A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services does not provide detailed monitoring with CloudWatch?

- A. AWS EMR
- B. AWS RDS
- C. AWS ELB
- D. AWS Route53

Answer: A

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. Services, such as RDS, EC2, Auto Scaling, ELB, and Route 53 can provide the monitoring data every minute.

NEW QUESTION 219

A user has configured ELB with SSL using a security policy for secure negotiation between the client and load balancer. Which of the below mentioned security policies is supported by ELB?

- A. Dynamic Security Policy
- B. All the other options
- C. Predefined Security Policy
- D. Default Security Policy

Answer: C

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. ELB supports two policies: Predefined Security Policy, which comes with predefined cipher and SSL protocols; Custom Security Policy, which allows the user to configure a policy.

NEW QUESTION 222

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 subnet
- B. Change the size of the subnets to 20.0.0.0/32 and 20.0.1.0/32, respectively
- 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 226

A user has launched 5 instances in EC2-CLASSIC and attached 5 elastic IPs to the five different instances in the US East region. The user is creating a VPC in the same region. The user wants to assign an elastic IP to the VPC instance. How can the user achieve this?

- A. The user has to request AWS to increase the number of elastic IPs associated with the account
- B. AWS allows 10 EC2 Classic IPs per region; so it will allow to allocate new Elastic IPs to the same region
- C. The AWS will not allow to create a new elastic IP in VPC; it will throw an error
- D. The user can allocate a new IP address in VPC as it has a different limit than EC2

Answer: D

Explanation: Section: (none)

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. A user can have 5 IP addresses per region with EC2 Classic. The user can have 5 separate IPs with VPC in the same region as it has a separate limit than EC2 Classic.

NEW QUESTION 231

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

- A. Data is automatically deleted
- B. Data is automatically saved as an EBS snapshot
- C. Data is unavailable until the instance is restarted
- D. Data is automatically saved as an EBS volume

Answer: A

NEW QUESTION 233

A user has setup a VPC with CIDR 20.0.0.0/16. The VPC has a private subnet (20.0.1.0/24) and a public subnet (20.0.0.0/24). The user's data centre has CIDR of 20.0.54.0/24 and 20.1.0.0/24. If the private subnet wants to communicate with the data centre, what will happen?

- A. It will allow traffic communication on both the CIDRs of the data centre
- B. It will not allow traffic with data centre on CIDR 20.1.0.0/24 but allows traffic communication on 20.0.54.0/24
- C. It will not allow traffic communication on any of the data centre CIDRs
- D. It will allow traffic with data centre on CIDR 20.1.0.0/24 but does not allow on 20.0.54.0/24

Answer: D

Explanation:

VPC allows the user to set up a connection between his VPC and corporate or home network data centre. If the user has an IP address prefix in the VPC that overlaps with one of the networks' prefixes, any traffic to the network's prefix is dropped. In this case CIDR 20.0.54.0/24 falls in the VPC's CIDR range of 20.0.0.0/16. Thus, it will not allow traffic on that IP. In the case of 20.1.0.0/24, it does not fall in the VPC's CIDR range. Thus, traffic will be allowed on it.

NEW QUESTION 238

A system admin is planning to encrypt all objects being uploaded to S3 from an application. The system admin does not want to implement his own encryption algorithm; instead he is planning to use server side encryption by supplying his own key (SSE-C). Which parameter is not required while making a call for SSE-C?

- A. x-amz-server-side-encryption-customer-key-AES-256
- B. x-amz-server-side-encryption-customer-key

- C. x-amz-server-side-encryption-customer-algorithm
- D. x-amz-server-side-encryption-customer-key-MD5

Answer: A

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.. When the user is supplying his own encryption key, the user has to send the below mentioned parameters as a part of the API calls: x-amz-server-side-encryption-customer-algorithm: Specifies the encryption algorithm x-amz-server-side-encryption-customer-key: To provide the base64-encoded encryption key x-amz-server-side-encryption-customer-key-MD5: To provide the base64-encoded 128-bit MD5 digest of the encryption key

NEW QUESTION 242

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 243

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 246

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 file
- 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 instance

Answer: BCE

NEW QUESTION 248

Which method can be used to prevent an IP address block from accessing public objects in an S3 bucket?

- A. Create a bucket policy and apply it to the bucket
- B. Create a NACL and attach it to the VPC of the bucket
- C. Create an ACL and apply it to all objects in the bucket
- D. Modify the IAM policies of any users that would access the bucket

Answer: A

Explanation: Reference:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/example-bucket-policies.html>

NEW QUESTION 253

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 257

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

- A. Allow Inbound on port 3306 for Source Web Server Security Group (WebSecGr)
- B. Allow Inbound on port 3306 from source 20.0.0.0/16
- C. Allow Outbound on port 3306 for Destination Web Server Security Group (WebSecGr)
- D. Allow Outbound on port 80 for Destination NAT Instance 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 to host the web server and DB server respectively, the user should configure that the instances in the private subnet can receive inbound traffic from the public subnet on the DB port. Thus, configure port 3306 in Inbound with the source as the Web Server Security Group (WebSecGrp). The user should configure ports 80 and 443 for Destination 0.0.0.0/0 as the route table directs traffic to the NAT instance from the private subnet.

NEW QUESTION 261

An organization has setup Auto Scaling with ELB. Due to some manual error, one of the instances got rebooted. Thus, it failed the Auto Scaling health check. Auto Scaling has marked it for replacement. How can the system admin ensure that the instance does not get terminated?

- A. Update the Auto Scaling group to ignore the instance reboot event
- B. It is not possible to change the status once it is marked for replacement
- C. Manually add that instance to the Auto Scaling group after reboot to avoid replacement
- D. Change the health of the instance to healthy using the Auto Scaling commands

Answer: D

Explanation:

After an instance has been marked unhealthy by Auto Scaling, as a result of an Amazon EC2 or ELB health check, it is almost immediately scheduled for replacement as it will never automatically recover its health. If the user knows that the instance is healthy then he can manually call the SetInstanceHealth action (or the as-setinstance-health command from CLI) to set the instance's health status back to healthy. Auto Scaling will throw an error if the instance is already terminating or else it will mark it healthy.

NEW QUESTION 262

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 263

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 267

A user wants to find the particular error that occurred on a certain date in the AWS MySQL RDS DB. Which of the below mentioned activities may help the user to get the data easily?

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

Answer: C

Explanation:

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

NEW QUESTION 269

Which services allow the customer to retain run administrative privileges or the underlying EC2 instances? Choose 2 answers

- A. AWS Elastic Beanstalk
- B. Amazon Elastic Map Reduce
- C. Elastic Load Balancing
- D. Amazon Relational Database Service
- E. Amazon Elasti Cache

Answer: AB

NEW QUESTION 272

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

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

Answer: B

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When an instance is launched it will have a network interface attached with it. The user cannot delete the subnet until he terminates the instance and deletes the network interface.

NEW QUESTION 273

A user has configured an EC2 instance in the US-East-1a zone. The user has enabled detailed monitoring of the instance. The user is trying to get the data from CloudWatch using a CLI. Which of the below mentioned CloudWatch endpoint URLs should the user use?

- A. monitoring.us-east-1.amazonaws.com
- B. monitoring.us-east-1-a.amazonaws.com
- C. monitoring.us-east-1a.amazonaws.com
- D. cloudwatch.us-east-1a.amazonaws.com

Answer: A

Explanation:

The CloudWatch resources are always region specific and they will have the end point as region specific. If the user is trying to access the metric in the US-East-1 region, the endpoint URL will be: monitoring.us-east-1.amazonaws.com

NEW QUESTION 278

How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

- A. Simply create a new volume in the other AZ and specify the original volume as the source
- B. Detach the volume, then use the ec2-migrate-volume command to move it to another AZ
- C. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ
- D. Detach the volume and attach it to another EC2 instance in the other AZ

Answer: D

Explanation: Reference:

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

NEW QUESTION 282

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 284

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 as --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 286

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 291

A user is observing the EC2 CPU utilization metric on CloudWatch. The user has observed some interesting patterns while filtering over the 1 week period for a particular hour. The user wants to zoom that data point to a more granular period. How can the user do that easily with CloudWatch?

- A. The user can zoom a particular period by selecting that period with the mouse and then releasing the mouse
- B. The user can zoom a particular period by double clicking on that period with the mouse
- C. The user can zoom a particular period by specifying the aggregation data for that period
- D. The user can zoom a particular period by specifying the period in the Time Range

Answer: A

NEW QUESTION 292

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 297

A sys admin has enabled a log on ELB. Which of the below mentioned activities are not captured by the log?

- A. Response processing time
- B. Front end processing time
- C. Backend processing time
- D. Request processing time

Answer: B

Explanation:

Elastic Load Balancing access logs capture detailed information for all the requests made to the load balancer. Each request will have details, such as client IP, request path, ELB IP, time, and latencies. The time will have information, such as Request Processing time, Backend Processing time and Response Processing time.

NEW QUESTION 299

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 300

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 303

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 308

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 310

How can you secure data at rest on an EBS volume?

- A. Encrypt the volume using the S3 server-side encryption service
- B. Attach the volume to an instance using EC2's SSL interface
- C. Create an IAM policy that restricts read and write access to the volume
- D. Write the data randomly instead of sequentially
- E. Use an encrypted file system on top of the EBS volume

Answer: C

Explanation:

Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/policies_examples.html

NEW QUESTION 311

A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services provides detailed monitoring with CloudWatch without charging the user extra?

- A. AWS Auto Scaling
- B. AWS Route 53
- C. AWS EMR
- D. AWS SNS

Answer: B

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. Services, such as RDS, ELB, OpsWorks, and Route 53 can provide the monitoring data every minute without charging the user.

NEW QUESTION 314

An application you maintain consists of multiple EC2 instances in a default tenancy VPC. This application has undergone an internal audit and has been determined to require dedicated hardware for one instance. Your compliance team has given you a week to move this instance to single-tenant hardware. Which process will have minimal impact on your application while complying with this requirement?

- A. Create a new VPC with `tenancy=dedicated` and migrate to the new VPC
- B. Use `ec2-reboot-instances` command line and set the parameter `"dedicated=true"`
- C. Right click on the instance, select properties and check the box for dedicated tenancy

D. Stop the instance, create an AMI, launch a new instance with tenancy=dedicated, and terminate the old instance

Answer: A

Explanation: Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/ApiReference-cmd-CreateVpc.html>

NEW QUESTION 318

A user has created a VPC with the public subnet. The user has created a security group for that VPC. Which of the below mentioned statements is true when a security group is created?

- A. It can connect to the AWS services, such as S3 and RDS by default
- B. It will have all the inbound traffic by default
- C. It will have all the outbound traffic by default
- D. It will by default allow traffic to the internet gateway

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. AWS provides two features the user can use to increase security in VPC: security groups and network ACLs. Security groups work at the instance level while ACLs work at the subnet level. When a user creates a security group with AWS VPC, by default it will allow all the outbound traffic but block all inbound traffic.

NEW QUESTION 320

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 324

You have private video content in S3 that you want to serve to subscribed users on the Internet. User IDs, credentials, and subscriptions are stored in an Amazon RDS database.

Which configuration will allow you to securely serve private content to your users?

- A. Generate pre-signed URLs for each user as they request access to protected S3 content
- B. Create an IAM user for each subscribed user and assign the GetObject permission to each IAM user
- C. Create an S3 bucket policy that limits access to your private content to only your subscribed users' credentials
- D. Create a CloudFront Origin Identity user for your subscribed users and assign the GetObject permission to this user

Answer: C

Explanation: Reference:

<https://java.awsblog.com/post/Tx1VE22EWFR4H86/Accessing-Private-Content-in-Amazon-CloudFront>

NEW QUESTION 328

A root account owner is trying to understand the S3 bucket ACL. Which of the below mentioned options cannot be used to grant ACL on the object using the authorized predefined group?

- A. Authenticated user group
- B. All users group
- C. Log Delivery Group
- D. Canonical user group

Answer: D

Explanation:

An S3 bucket ACL grantee can be an AWS account or one of the predefined Amazon S3 groups. Amazon S3 has a set of predefined groups. When granting account access to a group, the user can specify one of the URLs of that group instead of a canonical user ID. AWS S3 has the following predefined groups: Authenticated Users group: It represents all AWS accounts. All Users group: Access permission to this group allows anyone to access the resource. Log Delivery group: WRITE permission on a bucket enables this group to write server access logs to the bucket.

NEW QUESTION 329

An organization has created 10 IAM users. The organization wants each of the IAM users to have access to a separate DynamoDB table. All the users are added to the same group and the organization wants to setup a group level policy for this. How can the organization achieve this?

- A. Define the group policy and add a condition which allows the access based on the IAM name
- B. Create a DynamoDB table with the same name as the IAM user name and define the policy rule which grants access based on the DynamoDB ARN using a variable
- C. Create a separate DynamoDB database for each user and configure a policy in the group based on the DB variable
- D. It is not possible to have a group level policy which allows different IAM users to different DynamoDB Tables

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. AWS DynamoDB has only tables and the organization cannot make separate databases. The organization should create a table with the same name as the IAM user name and use the ARN of DynamoDB as part of the group policy. The sample policy is shown below:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": ["dynamodb:*"],
    "Resource": "arn:aws:dynamodb:region:account-number-without-hyphens:table/${aws:username}"
  }]
}
```

NEW QUESTION 331

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.e. 1×10^{126}).
- 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 332

Amazon EBS snapshots have which of the following two characteristics? (Choose 2.) Choose 2 answers

- A. EBS snapshots only save incremental changes from snapshot to snapshot
- B. EBS snapshots can be created in real-time without stopping an EC2 instance
- C. EBS snapshots can only be restored to an EBS volume of the same size or smaller
- D. EBS snapshots can only be restored and mounted to an instance in the same Availability Zone as the original EBS volume

Answer: AD

NEW QUESTION 336

An organization has launched 5 instances: 2 for production and 3 for testing. The organization wants that one particular group of IAM users should only access the test instances and not the production ones. How can the organization set that as a part of the policy?

- A. Launch the test and production instances in separate regions and allow region wise access to the group
- B. Define the IAM policy which allows access based on the instance ID
- C. Create an IAM policy with a condition which allows access to only small instances
- D. Define the tags on the test and production servers and add a condition to the IAM policy which allows access to specific tags

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 can add conditions as a part of the IAM policies. The condition can be set on AWS Tags, Time, and Client IP as well as on various parameters. If the organization wants the user to access only specific instances he should define proper tags and add to the IAM policy condition.

The sample policy is shown below.

```
"Statement": [
  {
    "Action": "ec2:*",
    "Effect": "Allow",
    "Resource": "*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/InstanceType": "Production"
      }
    }
  }
]
```

NEW QUESTION 341

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 and Glacier (Archiv
- E. 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 342

A sys admin is planning to subscribe to the RDS event notifications. For which of the below mentioned source categories the subscription cannot be configured?

- A. DB security group
- B. DB snapshot
- C. DB options group
- D. DB parameter group

Answer: C

Explanation:

Amazon RDS uses the Amazon Simple Notification Service (SNS) 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.

NEW QUESTION 346

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 348

A user is trying to send custom metrics to CloudWatch using the PutMetricData APIs. Which of the below mentioned points should the user needs to take care while sending the data to CloudWatch?

- A. The size of a request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests
- B. The size of a request is limited to 128KB for HTTP GET requests and 64KB for HTTP POST requests
- C. The size of a request is limited to 40KB for HTTP GET requests and 8KB for HTTP POST requests
- D. The size of a request is limited to 16KB for HTTP GET requests and 80KB for HTTP POST requests

Answer: A

Explanation:

With AWS CloudWatch, the user can publish data points for a metric that share not only the same time stamp, but also the same namespace and dimensions. CloudWatch can accept multiple data points in the same PutMetricData call with the same time stamp. The only thing that the user needs to take care of is that the size of a PutMetricData request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests.

NEW QUESTION 352

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 355

A user has launched an EC2 Windows instance from an instance store backed AMI. The user has also set the Instance initiated shutdown behavior to stop. What will happen when the user shuts down the OS?

- A. It will not allow the user to shutdown the OS when the shutdown behaviour is set to Stop
- B. It is not possible to set the termination behaviour to Stop for an Instance store backed AMI instance
- C. The instance will stay running but the OS will be shutdown
- D. The instance will be terminated

Answer: B

Explanation:

When the EC2 instance is launched from an instance store backed AMI, it will not allow the user to configure the shutdown behaviour to "Stop". It gives a warning that the instance does not have the EBS root volume.

NEW QUESTION 358

.....

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](#)