SOA-C01 Dumps

AWS Certified SysOps Administrator - Associate

https://www.certleader.com/SOA-C01-dumps.html

NEW QUESTION 1
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 2
You have started a new job and are reviewing your company's infrastructure on AWS You notice one web application where they have an Elastic Load Balancer (&B) in front of web instances in an Auto Scaling Group When you check the metrics for the ELB in CloudWatch you see four healthy instances in Availability Zone (AZ) A and zero in AZ B There are zero unhealthy instances.
What do you need to fix to balance the instances across AZs?

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

Answer: B

NEW QUESTION 3
You have identified network throughput as a bottleneck on your m1.small EC2 instance when uploading data Into Amazon S3 In the same region.
How do you remedy this situation?

A. Add an additional ENI
B. Change to a larger Instance
C. Use DirectConnect between EC2 and S3
D. Use EBS PIOPS on the local volume

Answer: B

Explanation:
Reference:
https://media.amazonwebservices.com/AWS_Amazon_EMR_Best_Practices.pdf

NEW QUESTION 4
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 balancer.
D. Create a second multi-AZ RDS instance in another Availability Zone and configure replication to provide a redundant database.

Answer: A

NEW QUESTION 5
The majority of your Infrastructure is on premises and you have a small footprint on AWS Your company has decided to roll out a new application that is heavily dependent on low latency connectivity to LOAP for authentication Your security policy requires minimal changes to the company's existing application user management processes.
What option would you implement to successfully launch this application1?

A. Create a second, independent LOAP server in AWS for your application to use for authentication
B. Establish a VPN connection so your applications can authenticate against your existing on- premises LDAP servers
C. Establish a VPN connection between your data center and AWS create a LDAP replica on AWS and configure your application to use the LDAP replica for authentication
D. Create a second LDAP domain on AWS establish a VPN connection to establish a trust relationship between your new and existing domains and use the new domain for authentication

Answer: C

Explanation:
Since it requires no changes to the authentication infrastructure as requested in the question. Option D creates a new LDAP, trusts, etc.

NEW QUESTION 6
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 7
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 ENI0 that can be moved to a failover instance

Answer: BD

Explanation:
This is an odd question. First of all, option A cannot be right because ELB does not failover. Cannot be C because Route 53 does work with hard coded IP. Only B & D cannot be rule out so best answer.

NEW QUESTION 8
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 grou
D. Specify the Auto Scaling group to include multiple AZs but have a min-size of 1 and max-size of 1
E. Configure an ELB in front of the bastion instance

Answer: C

NEW QUESTION 9
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 192 168 100 0 full access to the "mybucket" bucket
B. Denies the server with the IP address 192 168 100 188 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: B

NEW QUESTION 10
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 10
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 Doth 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 duster

Answer: AB

Explanation:
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheMetrics.WhichShouldIMo nitor.html

NEW QUESTION 11
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 lo fire an alarm, and the alarm notifies your operations team by email and SMS of a possible application health problem. However, you also need to ??watch the watcher?? --the monitoring instance itself - and be notified if it becomes unhealthy.
Which of the following is a simple way to achieve that goal?

A. Run another monitoring instance that pings the monitoring instance and fires a could watch alarm mat 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 nave the alarm notify your operations team if the CPU usage exceeds 50% few more than one minute; then have your monitoring application go into a CPU-bound loop should it Detect any application problems.
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 (he role of the previous monitoring instance and beginning adding messages to the SQS queue.

Answer: B

NEW QUESTION 14
You are attempting to connect to an instance in Amazon VPC without success. You have already verified that the VPC has an Internet Gateway (IGW) the instance has an associated Elastic IP (EIP) and correct security group rules are in place.
Which VPC component should you evaluate next?

A. The configuration of a NAT instance
B. The configuration of the Routing Table
C. The configuration of the internet Gateway (IGW)
D. The configuration of SRC/DST checking

Answer: B

Explanation:
Reference:
http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/UserScenariosForVPC.ht ml

NEW QUESTION 18
You are tasked with the migration of a highly trafficked Node JS application to AWS In order to comply with organizational standards Chef recipes must be used to configure the application servers that host this application and to support application lifecycle events.
Which deployment option meets these requirements while minimizing administrative burden?

A. Create a new stack within Opsworks add the appropriate layers to the stack and deploy the application
B. Create a new application within Elastic Beanstalk and deploy this application to a new environment
C. Launch a Node.JS server from a community AMI and manually deploy the application to the launched EC2 instance
D. Launch and configure Chef Server on an EC2 instance and leverage the AWS CLI to launch application servers and configure those instances using Chef.

Answer: A

Explanation:
OpsWorks has integrated support for Chef and lifecycle events.
See: http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook.html

NEW QUESTION 22
What are characteristics of Amazon S3? Choose 2 answers

A. Objects are directly accessible via a URL
B. S3 should be used to host a relational database
C. S3 allows you to store objects or virtually unlimited size
D. S3 allows you to store virtually unlimited amounts of data
E. S3 offers Provisioned IOPS

Answer: AD

NEW QUESTION 23
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.BackingUpAndRestoringAma zonRDSInstances.html

NEW QUESTION 27
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

Explanation:
It seems that this question wants to emphasize below (S3 Faq ?V https://aws.amazon.com/s3/faqs/) You specify a region when you create your Amazon S3 bucket. Within that region, your objects are redundantly stored on multiple devices across multiple facilities. Please refer to Regional Products and Services for details of Amazon S3 service availability by region.

NEW QUESTION 28
You have a server with a 5O0GB Amazon EBS data volume. The volume is 80% full. You need to back up the volume at regular intervals and be able to re-create the volume in a new Availability Zone in the shortest time possible. All applications using the volume can be paused for a period of a few minutes with no discernible user impact.
Which of the following backup methods will best fulfill your requirements?

A. Take periodic snapshots of the EBS volume
B. Use a third party Incremental backup application to back up to Amazon Glacier
C. Periodically back up all data to a single compressed archive and archive to Amazon S3 using a parallelized multi-part upload
D. Create another EBS volume in the second Availability Zone attach it to the Amazon EC2 instance, and use a disk manager to mirror me two disks

Answer: A

Explanation:
Since an EBS volume should be in the same AZ as the EC2 instance. You cannot connect a EBS volume in another AZ.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-restoring-volume.html EBS volumes can only be attached to EC2 instances within the same Availability Zone.

NEW QUESTION 33
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 oft of your web application tier
D. Use EBS PIOPs to serve the image faster out of your EC2 instances

Answer: AB

Explanation:
CloudFront gets the image closer to the user and Route53 ensures the best connection based on network latency. Option D does not address the issue.

NEW QUESTION 35
If you want to launch Amazon Elastic Compute Cloud (EC2) Instances and assign each Instance a predetermined private IP address you should:

A. Assign a group or sequential Elastic IP address to the instances

B. Launch the instances in a Placement Group
C. Launch the instances in the Amazon virtual Private Cloud (VPC).
D. Use standard EC2 instances since each instance gets a private Domain Name Service (DNS) already
E. Launch the Instance from a private Amazon Machine image (Mil)

Answer: C

Explanation:
When you launch an instance into a VPC, a primary private IP address from the address range of the subnet is assigned to the default network interface (eth0) of the instance. If you don??t specify a primary private IP address, we select an available IP address in the subnet range for you
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-ip-addressing.html

NEW QUESTION 37
What would happen to an RDS (Relational Database Service) multi-Availability Zone deployment of the primary OB instance fails?

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

Answer: D

Explanation:
https://aws.amazon.com/rds/faqs/

NEW QUESTION 41
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 43
A user has developed an application which is required to send the data to a NoSQL database. The user wants to decouple the data sending such that the application keeps processing and sending data but does not wait for an acknowledgement of DB. Which of the below mentioned applications helps in this scenario?

A. AWS Simple Notification Service
B. AWS Simple Workflow
C. AWS Simple Queue Service
D. AWS Simple Query Service

Answer: C

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. In this case, the user can use AWS SQS to send messages which are received from an application and sent to DB. The application can continue processing data without waiting for any acknowledgement from DB. The user can use SQS to transmit any volume of data without losing messages or requiring other services to always be available.

NEW QUESTION 44
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 capacit
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 47
An organization is planning to create 5 different AWS accounts considering various security requirements. The organization wants to use a single payee account by using the consolidated billing option. Which of the below mentioned statements is true with respect to the above information?

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

Answer: B

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

NEW QUESTION 51
A user has deployed an application on his private cloud. The user is using his own monitoring tool. He wants to configure that whenever there is an error, the monitoring tool should notify him via SMS. Which of the below mentioned AWS services will help in this scenario?

A. None because the user infrastructure is in the private cloud
B. AWS SNS
C. AWS SES
D. AWS SMS

Answer: B

Explanation:
Amazon Simple Notification Service (Amazon SNS. is a fast, flexible, and fully managed push messaging service. Amazon SNS can be used to make push notifications to mobile devices. Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service
(SQS. queues or to any HTTP endpoint. In this case user can use the SNS apis to send SMS.

NEW QUESTION 55
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.
http://docs.aws.amazon.com/cli/latest/reference/opsworks/set-time-based-auto-scaling.html

NEW QUESTION 58
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.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/AlarmThatSendsEmail.ht ml

NEW QUESTION 63
An organization is setting up programmatic billing access for their AWS account. Which of the below mentioned services is not required or enabled when the organization wants to use programmatic access?

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

Answer: C

Explanation:
AWS provides an option to have programmatic access to billing. Programmatic Billing Access leverages the existing Amazon Simple Storage Service (Amazon S3.

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

NEW QUESTION 68
An admin is planning to monitor the ELB. Which of the below mentioned services does not help the admin capture the monitoring information about the ELB activity?

A. ELB Access logs
B. ELB health check
C. CloudWatch metrics
D. ELB API calls with CloudTrail

Answer: B

Explanation:
The admin can capture information about Elastic Load Balancer using either:
CloudWatch Metrics ELB Logs files which are stored in the S3 bucket CloudTrail with API calls which can notify the user as well generate logs for each API calls
The health check is internally performed by ELB and does not help the admin get the ELB activity.

NEW QUESTION 69
A user is planning to use AWS CloudFormation. Which of the below mentioned functionalities does not help him to correctly understand CloudFormation?

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

Answer: A

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

NEW QUESTION 73
A user has launched 10 instances from the same AMI ID using Auto Scaling. The user is trying to see the average CPU utilization across all instances of the last 2 weeks under the CloudWatch console. How can the user achieve this?

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

Answer: A

Explanation:
Auto Scaling has its own aggregated CPU Utilization metric.

NEW QUESTION 75
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 the following transports as part of the subscription requests: ??HTTP??, ??HTTPS??,??Email??, ??Email-JSON??, ??SQS??, ??and SMS??.

NEW QUESTION 80
A user has configured an Auto Scaling group with ELB. The user has enabled detailed CloudWatch monitoring on Auto Scaling. Which of the below mentioned statements will help the user understand the functionality better?

A. It is not possible to setup detailed monitoring for Auto Scaling
B. In this case, Auto Scaling will send data every minute and will charge the user extra
C. Detailed monitoring will send data every minute without additional charges
D. Auto Scaling sends data every minute only and does not charge the user

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html CloudWatch monitors the following services. As soon as you begin using a service, it automatically sends metrics to CloudWatch for you.

CloudWatch offers either basic or detailed monitoring for supported AWS products. Basic monitoring means that a service sends data points to CloudWatch every five minutes. Detailed monitoring means that a service sends data points to CloudWatch every minute.

Note

If you are using a service that supports both basic and detailed data collection (for example, Amazon EC2 and Auto Scaling), and you want to access detailed statistics, you must enable detailed metric collection for that service.

Auto Scaling

Auto Scaling sends data to CloudWatch every 5 minutes by default. For an additional charge, you can enable detailed monitoring for Auto Scaling, which sends data to CloudWatch every minute. You can create alarms using Auto Scaling Dimensions and Metrics. For more information, see Monitor Your

Auto Scaling Instances in the Auto Scaling User Guide.

NEW QUESTION 84
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 SQSdoes 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 89
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 91
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 95
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 100
An organization has created 50 IAM users. The organization wants that each user can change their password but cannot change their access keys. How can the organization achieve this?

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

Answer: D

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

NEW QUESTION 103
An application is generating a log file every 5 minutes. The log file is not critical but may be required only for verification in case of some major issue. The file should be accessible over the internet whenever required. Which of the below mentioned options is a best possible storage solution for it?

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

Answer: D

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

NEW QUESTION 106
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 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.

NEW QUESTION 110
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 grou

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 111
A user has configured ELB with three instances. The user wants to achieve High Availability as well as redundancy with ELB. Which of the below mentioned AWS services helps the user achieve this for ELB?

A. Route 53
B. AWS Mechanical Turk
C. Auto Scaling
D. AWS EMR

Answer: A

Explanation:

The user can provide high availability and redundancy for applications running behind Elastic Load Balancer by enabling the Amazon Route 53 Domain Name System (DNS. failover for the load balancers. Amazon Route 53 is a DNS service that provides reliable routing to the user??s infrastructure.

NEW QUESTION 115
An organization is using AWS since a few months. The finance team wants to visualize the pattern of AWS spending. Which of the below AWS tool will help for this requirement?

A. AWS Cost Manager
B. AWS Cost Explorer
C. AWS CloudWatch
D. AWS Consolidated Billing

Answer: B

Explanation:
The AWS Billing and Cost Management console includes the Cost Explorer tool for viewing AWS cost data as a graph. It does not charge extra to user for this service. With Cost Explorer the user can filter graphs using resource tags or with services in AWS. If the organization is using Consolidated Billing it helps generate report based on linked accounts. This will help organization to identify areas that require further inquiry. The organization can view trends and use that to understand spend and to predict future costs.

NEW QUESTION 117
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 is planning to setup notifications on the RDS DB for a snapshot. Which of the below mentioned event categories is not supported by RDS for this snapshot source type?

A. Backup
B. Creation
C. Deletion
D. Restoration

Answer: A

Explanation:
Amazon RDS uses the Amazon Simple Notification Service to provide a notification when an Amazon RDS event occurs. Event categories for a snapshot source type include: Creation, Deletion, and Restoration. The Backup is a part of DB instance source type.

NEW QUESTION 121
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 124
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 manuall
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. I t 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 126
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 130
A user has stored data on an encrypted EBS volume. The user wants to share the data with his friend??s AWS account. How can user achieve this?

A. Create an AMI from the volume and share the AMI
B. Copy the data to an unencrypted volume and then share
C. Take a snapshot and share the snapshot with a friend
D. If both the accounts are using the same encryption key then the user can share the volume directly

Answer: B

Explanation:
AWS EBS supports encryption of the volume. It also supports creating volumes from existing snapshots
provided the snapshots are created from encrypted volumes. If the user is having data on an encrypted volume and is trying to share it with others, he has to copy the data from the encrypted volume to a new unencrypted volume. Only then can the user share it as an encrypted volume data. Otherwise the snapshot cannot be shared.

NEW QUESTION 132
A user has created a VPC with CIDR 20.0.0.0/16 with only a private subnet and VPN connection using the VPC wizard. The user wants to connect to the instance in a private subnet over SSH. How should the user define the security rule for SSH?

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

Answer: A

Explanation:
The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data center, the user can setup a case with a VPN only subnet (private. which uses VPN access to connect with his data center. When the user has configured this setup with Wizard, all network connections to the instances in the subnet will come from his data center. The user has to configure the security group of the private subnet which allows the inbound traffic on SSH (port 22. from the data center??s network range.

NEW QUESTION 136
A user has configured ELB with two EBS backed EC2 instances. The user is trying to understand the DNS access and IP support for ELB. Which of the below mentioned statements may not help the user understand the IP mechanism supported by ELB?

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

Answer: D

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

NEW QUESTION 137
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 139
A user is checking the CloudWatch metrics from the AWS console. The user notices that the CloudWatch data is coming in UTC. The user wants to convert the data to a local time zone. How can the user perform this?

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

Answer: B

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

NEW QUESTION 144
An organization (Account ID 123412341234. has attached the below mentioned IAM policy to a user. What does this policy statement entitle the user to perform?
"Statement": [
{
"Sid": "AllowUsersAllActionsForCredentials", "Effect": "Allow",
"Action": [
"iam:*AccessKey*",
],
"Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
}
]

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

Answer: D

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (Account ID 123412341234. wants some of their users to manage keys (access and secret access keys. of all IAM users, the organization should set the below mentioned policy which entitles the IAM user to modify keys of all IAM users with CLI, SDK or API.
"Statement": [
{
"Sid": "AllowUsersAllActionsForCredentials", "Effect": "Allow",
"Action": ["iam:*AccessKey*",
],
"Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
}
]

NEW QUESTION 147
A user has launched an EBS backed EC2 instance. What will be the difference while performing the
restart or stop/start options on that instance?

A. For restart it does not charge for an extra hour, while every stop/start it will be charged as a separate hour
B. Every restart is charged by AWS as a separate hour, while multiple start/stop actions during a single hour will be counted as a single hour
C. For every restart or start/stop it will be charged as a separate hour
D. For restart it charges extra only once, while for every stop/start it will be charged as a separate hour

Answer: A

Explanation:
For an EC2 instance launched with an EBS backed AMI, each time the instance state is changed from stop to start/ running, AWS charges a full instance hour, even if these transitions happen multiple times within a single hour. Anyway, rebooting an instance AWS does not charge a new instance billing hour.

NEW QUESTION 149
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
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 151
A sysadmin 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 terminat instances, respectively. Some other actions performed by Auto Scaling are: AddToLoadbalancer, AlarmNotification, HealthCheck, AZRebalance, ReplaceUnHealthy, and ScheduledActions.

NEW QUESTION 155
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 158
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 162
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 166
A user has created a VPC with CIDR 20.0.0.0/24. The user has created a public subnet with CIDR 20.0.0.0/25 and a private subnet with CIDR 20.0.0.128/25. The user has launched one instance each in the private and public subnets. Which of the below mentioned options cannot be the correct IP address (private IP. assigned to an instance in the public or private subnet?

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

Answer: A

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

NEW QUESTION 167
A user has launched an EBS backed EC2 instance. The user has rebooted the instance. Which of the below mentioned statements is not true with respect to the reboot action?

A. The private and public address remains the same
B. The Elastic IP remains associated with the instance
C. The volume is preserved
D. The instance runs on a new host computer

Answer: D

Explanation:
A user can reboot an EC2 instance using the AWS console, the Amazon EC2 CLI or the Amazon EC2 API. Rebooting an instance is equivalent to rebooting an operating system. However, it is recommended that the user use the Amazon EC2 to reboot the instance instead of running the operating system reboot command from the instance. The instance remains on the same host computer and maintains its public DNS name, private IP address, and any data on its instance store volumes. It typically takes a few minutes for the reboot to complete, but the time it takes to reboot depends on the instance configuration.

NEW QUESTION 168
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 similartime 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 169
A user is planning to schedule a backup for an EBS volume. The user wants security of the snapshot data. How can the user achieve data encryption with a snapshot?

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

Answer: A

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

NEW QUESTION 170
A user has setup an EBS backed instance and attached 2 EBS volumes to it. The user has setup a CloudWatch alarm on each volume for the disk data. The user has stopped the EC2 instance and detached the EBS volumes. What will be the status of the alarms on the EBS volume?

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

Answer:

B

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

NEW QUESTION 172
An organization is measuring the latency of an application every minute and storing data inside a file in the JSON format. The organization wants to send all latency data to AWS CloudWatch. How can the organization achieve this?

A. The user has to parse the file before uploading data to CloudWatch
B. It is not possible to upload the custom data to CloudWatch
C. The user can supply the file as an input to the CloudWatch command
D. The user can use the CloudWatch Import command to import data from the file to CloudWatch

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. The user has to always include the namespace as part of the request. If the user wants to upload the custom data from a file, he can supply file name along with the parameter -- metric-data to command put-metric-data.

NEW QUESTION 173
A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console?

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

Answer: D

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

NEW QUESTION 176
A user has launched multiple EC2 instances for the purpose of development and testing in the same region. The user wants to find the separate cost for the production and development instances. How can the user find the cost distribution?

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

Answer: D

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

NEW QUESTION 181
A user is planning to set up the Multi AZ feature of RDS. Which of the below mentioned conditions won't take advantage of the Multi AZ feature?

A. Availability zone outage
B. A manual failover of the DB instance using Reboot with failover option
C. Region outage
D. When the user changes the DB instance??s server type

Answer: C

Explanation:
Amazon RDS when enabled with Multi AZ will handle failovers automatically. Thus, the user can resume database operations as quickly as possible without administrative intervention. The primary DB instance switches over automatically to the standby replica if any of the following conditions occur:
An Availability Zone outage The primary DB instance fails
The DB instance's server type is changed
The DB instance is undergoing software patching
A manual failover of the DB instance was initiated using Reboot with failover

NEW QUESTION 183

An organization (account ID 123412341234. has configured the IAM policy to allow the user to modify his credentials. What will the below mentioned statement allow the user to perform?
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
],
"Resource": "arn:aws:iam:: 123412341234:group/TestingGroup"
}]

A. The IAM policy will throw an error due to an invalid resource name
B. The IAM policy will allow the user to subscribe to any IAM group
C. Allow the IAM user to update the membership of the group called TestingGroup
D. Allow the IAM user to delete the TestingGroup

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 organization (account ID 123412341234. wants their users to manage their subscription to the groups, they should create a relevant policy for that. The below mentioned policy allows the respective IAM user to update the membership of the group called MarketingGroup.
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
],
"Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup "
}]

NEW QUESTION 184
A user has configured ELB with two EBS backed instances. The user has stopped the instances for 1 week to save costs. The user restarts the instances after 1 week. Which of the below mentioned statements will help the user to understand the ELB and instance registration better?

A. There is no way to register the stopped instances with ELB
B. The user cannot stop the instances if they are registered with ELB
C. If the instances have the same Elastic IP assigned after reboot they will be registered with ELB
D. The instances will automatically get registered with ELB

Answer: C

Explanation:
Elastic Load Balancing registers the user??s load balancer with his EC2 instance using the associated IP address. When the instances are stopped and started back they will have a different IP address. Thus, they will not get registered with ELB unless the user manually registers them. If the instances are assigned the same Elastic IP after reboot they will automatically get registered with ELB.

NEW QUESTION 185
A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered. What time out period should the user specify for connection draining?

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

Answer: B

NEW QUESTION 188
A user is using the AWS EC2. The user wants to make so that when there is an issue in the EC2 server, such as instance status failed, it should start a new instance in the user??s private cloud. Which AWS service helps to achieve this automation?

A. AWS CloudWatch + Cloudformation
B. AWS CloudWatch + AWS AutoScaling + AWS ELB
C. AWS CloudWatch + AWS VPC
D. AWS CloudWatch + AWS SNS

Answer: D

Explanation:
Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS. queues or to any HTTP endpoint. The user can configure a web service (HTTP End point. in his data centre which receives data and launches an instance in the private cloud. The user should configure the CloudWatch alarm to send a notification to SNS when the ??StatusCheckFailed?? metric is true for the EC2 instance. The SNS topic can be configured to send a notification to the user??s HTTP end point which launches an instance in the private cloud.

NEW QUESTION 192
An organization has created one IAM user and applied the below mentioned policy to the user. What entitlements do the IAM users avail with this policy?
{

```
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "ec2:Describe*", "Resource": "*"
},
{
"Effect": "Allow"
"Action": [ "cloudwatch:ListMetrics", "cloudwatch:GetMetricStatistics", "cloudwatch:Describe*"
],
"Resource": "*"
},
{
"Effect": "Allow",
"Action": "autoscaling:Describe*", "Resource": "*"
}
]
}
```

A. The policy will allow the user to perform all read only activities on the EC2 services
B. The policy will allow the user to list all the EC2 resources except EBS
C. The policy will allow the user to perform all read and write activities on the EC2 services
D. The policy will allow the user to perform all read only activities on the EC2 services except load Balancing

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. If an organization wants to setup read only access to EC2 for a particular user, they should mention the action in the IAM policy which entitles the user for Describe rights for EC2, CloudWatch, Auto Scaling and ELB. In the policy shown below, the user will have read only access for EC2 and EBS, CloudWatch and Auto Scaling. Since ELB is not mentioned as a
part of the list, the user will not have access to ELB.

```
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "ec2:Describe*", "Resource": "*"
},
{
"Effect": "Allow", "Action": [ "cloudwatch:ListMetrics",
"cloudwatch:GetMetricStatistics", "cloudwatch:Describe*"
],
"Resource": "*"
},
{
"Effect": "Allow",
"Action": "autoscaling:Describe*", "Resource": "*"
}
]
}
```

NEW QUESTION 194
A user has enabled session stickiness with ELB. The user does not want ELB to manage the cookie; instead he wants the application to manage the cookie. What will happen when the server instance, which is bound to a cookie, crashes?

A. The response will have a cookie but stickiness will be deleted
B. The session will not be sticky until a new cookie is inserted
C. ELB will throw an error due to cookie unavailability
D. The session will be sticky and ELB will route requests to another server as ELB keeps replicating the Cookie

Answer: B

Explanation:
With Elastic Load Balancer, if the admin has enabled a sticky session with application controlled stickiness, the load balancer uses a special cookie generated by the application to associate the session with the original server which handles the request. ELB follows the lifetime of the application-generated cookie corresponding to the cookie name specified in the ELB policy configuration. The load balancer only inserts a new stickiness cookie if the application response includes a new application cookie. The load balancer stickiness cookie does not update with each request. If the application cookie is explicitly removed or expires, the session stops being sticky until a new application cookie is issued.

NEW QUESTION 196
A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. Which of the below mentioned statements is true with respect to this scenario?

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

Answer: D

Explanation:
A Virtual Private Cloud (VPC. is a virtual network dedicated to the user??s AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP then it will only have a private IP when launched. The user cannot connect to the instance from the internet. If the user wants an elastic IP to connect to the instance from the internet he should create an internet gateway and assign an elastic IP to instance.

NEW QUESTION 197
An organization has applied the below mentioned policy on an IAM group which has selected the IAM users. What entitlements do the IAM users avail with this policy?
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "*",
"Resource": "*"
}
]
}

A. The policy is not created correctl
B. It will throw an error for wrong resource name
C. The policy is for the grou
D. Thus, the IAM user cannot have any entitlement to this
E. It allows full access to all AWS services for the IAM users who are a part of this group
F. If this policy is applied to the EC2 resource, the users of the group will have full access to the EC2 Resources

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. The IAM group allows the organization to specify permissions for a collection of users. With the below mentioned policy, it will allow the group full access (Admin. to all AWS services.
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "*",
"Resource": "*"
}
]
}

NEW QUESTION 200
A user is configuring a CloudWatch alarm on RDS to receive a notification when the CPU utilization of RDS is higher than 50%. The user has setup an alarm when there is some inactivity on RDS, such as RDS unavailability. How can the user configure this?

A. Setup the notification when the CPU is more than 75% on RDS
B. Setup the notification when the state is Insufficient Data
C. Setup the notification when the CPU utilization is less than 10%
D. It is not possible to setup the alarm on RDS

Answer: B

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 alarm has three states: Alarm, OK and Insufficient data. The Alarm will change to Insufficient Data when any of the three situations arise: when the alarm has just started, when the metric is not available or when enough data is not available for the metric to determine the alarm state. If the user wants to find that RDS is not available, he can setup to receive the notification when the state is in Insufficient data.

NEW QUESTION 203
A user has created a VPC with a subnet and a security group. The user has launched an instance in that subnet and attached a public IP. The user is still unable to connect to the instance. The internet gateway has also been created. What can be the reason for the error?

A. The internet gateway is not configured with the route table
B. The private IP is not present
C. The outbound traffic on the security group is disabled
D. The internet gateway is not configured with the security group

Answer: A

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. When a user launches an instance and wants to connect to an instance, he needs an internet gateway. The internet gateway should be configured with the route table to allow traffic from the internet.

NEW QUESTION 204

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 208
George has launched three EC2 instances inside the US-East-1a zone with his AWS account. Ray has launched two EC2 instances in the US-East-1a zone with his AWS account. Which of the below entioned statements will help George and Ray understand the availability zone (AZ. concept better?

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

Answer: D

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

NEW QUESTION 209
A user has setup a CloudWatch alarm on the EC2 instance for CPU utilization. The user has setup to receive a notification on email when the CPU utilization is higher than 60%. The user is running a virus scan on the same instance at a particular time. The user wants to avoid receiving an email at this time. What should the user do?

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

Answer: B

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

NEW QUESTION 210
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 (WebSecGrp.
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 (WebSecGrp.
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 214
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 218
A user has launched an RDS MySQL DB with the Multi AZ feature. The user has scheduled the scaling of instance storage during maintenance window. What is the correct order of events during maintenance window?
Perform maintenance on standby Promote standby to primary
Perform maintenance on original primary Promote original master back as primary

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

Answer: B

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

NEW QUESTION 219
A sys admin is using server side encryption with AWS S3. Which of the below mentioned statements helps the user understand the S3 encryption functionality?

A. The server side encryption with the user supplied key works when versioning is enabled
B. The user can use the AWS console, SDK and APIs to encrypt or decrypt the content for server side encryption with the user supplied key
C. The user must send an AES-128 encrypted key
D. The user can upload his own encryption key to the S3 console

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. The encryption with the user supplied key (SSE-C. does not work with the AWS console. The S3 does not store the keys and the user has to send a key with each request. The SSE-C works when the user has enabled versioning.

NEW QUESTION 221
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 225
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 226
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 231
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 232
A user is receiving a notification from the RDS DB whenever there is a change in the DB security group. The user does not want to receive these notifications for only a month. Thus, he does not want to delete the notification. How can the user configure this?

A. Change the Disable button for notification to ??Yes?? in the RDS console
B. Set the send mail flag to false in the DB event notification console
C. The only option is to delete the notification from the console
D. Change the Enable button for notification to ??No?? in the RDS console

Answer: D

Explanation:
Amazon RDS uses the Amazon Simple Notification Service to provide a notification when an Amazon RDS event occurs. Event notifications are sent to the addresses that the user has provided while creating the subscription. The user can easily turn off the notification without deleting a subscription by setting the Enabled radio button to No in the Amazon RDS console or by setting the Enabled parameter to false using the CLI or Amazon RDS API.

NEW QUESTION 234
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 by mistake. The user is trying to create another subnet of CIDR 20.0.0.1/24. How can the user create the second subnet?

A. There is no need to update the subnet as VPC automatically adjusts the CIDR of the first subnet based on the second subnet??s CIDR
B. The user can modify the first subnet CIDR from the console
C. It is not possible to create a second subnet as one subnet with the same CIDR as the VPC has been created
D. The user can modify the first subnet CIDR with AWS CLI

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 the 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. The user cannot modify the CIDR of a subnet once it is created. Thus, in this case if required, the user has to delete the subnet and create new subnets.

NEW QUESTION 236
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 241
A user is using the AWS SQS to decouple the services. Which of the below mentioned operations is not supported by SQS?

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

Answer: D

Explanation:
Amazon Simple Queue Service (SQS. is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. The user can perform the following set of operations using the Amazon SQS: CreateQueue, ListQueues, DeleteQueue, SendMessage, SendMessageBatch, ReceiveMessage, DeleteMessage, DeleteMessageBatch, ChangeMessageVisibility, ChangeMessageVisibilityBatch, SetQueueAttributes, GetQueueAttributes, GetQueueUrl, AddPermission and RemovePermission. Operations can be performed only by the AWS account owner or an AWS account that the account owner has delegated to.

NEW QUESTION 245
A user has launched an EC2 instance from an instance store backed AMI. If the user restarts the instance, what will happen to the ephermal storage data?

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

Answer: D

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

NEW QUESTION 250
A user has configured an ELB to distribute the traffic among multiple instances. The user instances are facing some issues due to the back-end servers. Which of the below mentioned CloudWatch
metrics helps the user understand the issue with the instances?

A. HTTPCode_Backend_3XX
B. HTTPCode_Backend_4XX
C. HTTPCode_Backend_2XX
D. HTTPCode_Backend_5XX

Answer: D

Explanation:
CloudWatch is used to monitor AWS as well as the custom services. For ELB, CloudWatch provides various metrics including error code by ELB as well as by back-end servers (instances.. It gives data for the count of the number of HTTP response codes generated by the back-end instances. This metric does not include any response codes generated by the load balancer. These metrics are:
The 2XX class status codes represents successful actions
The 3XX class status code indicates that the user agent requires action The 4XX class status code represents client errors
The 5XX class status code represents back-end server errors

NEW QUESTION 253
A user has created a VPC with a public subnet. The user has terminated all the instances which are part of the subnet. Which of the below mentioned statements is true with respect to this scenario?

A. The user cannot delete the VPC since the subnet is not deleted
B. All network interface attached with the instances will be deleted
C. When the user launches a new instance it cannot use the same subnet
D. The subnet to which the instances were launched with will be deleted

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. When the user terminates the instance all the network interfaces attached with it are also deleted.

NEW QUESTION 254
A user has created a VPC with public and private subnets. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.1.0/24 and the public subnet uses CIDR 20.0.0.0/24. 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 of the NAT instance. Which of the below mentioned entries is not required for the NAT security group?

A. For Inbound allow Source: 20.0.1.0/24 on port 80
B. For Outbound allow Destination: 0.0.0.0/0 on port 80
C. For Inbound allow Source: 20.0.0.0/24 on port 80
D. For Outbound allow Destination: 0.0.0.0/0 on port 443

Answer: C

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 connect to the internet using the NAT instances. The user should first configure that NAT can receive traffic on ports 80 and 443 from the private subnet. Thus, allow ports 80 and 443 in Inbound for the private subnet 20.0.1.0/24. Now to route this traffic to the internet configure ports 80 and 443 in Outbound with destination 0.0.0.0/0. The NAT should not have an entry for the public subnet CIDR.

NEW QUESTION 259
A user has created an application which will be hosted on EC2. The application makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK to connect with from the EC2 instance. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

A. The user should attach an IAM role with DynamoDB access to the EC2 instance
B. The user should create an IAM user with DynamoDB access and use its credentials within the application to connect with DynamoDB
C. The user should create an IAM role, which has EC2 access so that it will allow deploying the application
D. The user should create an IAM user with DynamoDB and EC2 acces
E. Attach the user with the application so that it does not use the root account credentials

Answer: A

Explanation:
With AWS IAM a user is creating an application which runs on an EC2 instance and makes requests to AWS, such as DynamoDB or S3 calls. Here it is recommended that the user should not create an IAM user and pass the user's credentials to the application or embed those credentials inside the application. Instead, the user should use roles for EC2 and give that role access to DynamoDB /S3. When the roles are attached to EC2, it will give temporary security credentials to the application hosted on that EC2, to connect with DynamoDB / S3.

NEW QUESTION 260
An organization (Account ID 123412341234. has attached the below mentioned IAM policy to a user. What does this policy statement entitle the user to perform?
{
"Version": "2012-10-17",
"Statement": [{
"Sid": "AllowUsersAllActionsForCredentials", "Effect": "Allow",
"Action": ["iam:*LoginProfile", "iam:*AccessKey*",
"iam:*SigningCertificate*"
],
"Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
}]
}

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

Answer: D

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (Account ID 123412341234. wants some of their users to manage credentials (access keys, password, and sing in certificates. of all IAM users, they should set an applicable policy to that user or group of users. The below mentioned policy allows the IAM user to modify the credentials of all IAM user??s using only CLI, SDK or APIs. The user cannot use the AWS console for this activity since he does not have list permission for the IAM users.
{
"Version": "2012-10-17",
"Statement": [{
"Sid": "AllowUsersAllActionsForCredentials", "Effect": "Allow"
"Action": ["iam:*LoginProfile", "iam:*AccessKey*", "iam:*SigningCertificate*"
],
"Resource": ["arn:aws:iam::123412341234:user/${aws:username}"]
}]
}

NEW QUESTION 263
A user has a weighing plant. The user measures the weight of some goods every 5 minutes and sends data to AWS CloudWatch for monitoring and tracking. Which of the below mentioned parameters is mandatory for the user to include in the request list?

A. Value
B. Namespace
C. Metric Name
D. Timezone

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

NEW QUESTION 268
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 270
A user has provisioned 2000 IOPS to the EBS volume. The application hosted on that EBS is experiencing less IOPS than provisioned. Which of the below mentioned options does not affect the IOPS of the volume?

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

Answer: D

Explanation:
When the application does not experience the expected IOPS or throughput of the PIOPS EBS volume that was provisioned, the possible root cause could be that the EC2 bandwidth is the limiting factor and the instance might not be either EBS-optimized or might not have 10 Gigabit network connectivity. Another possible cause for not experiencing the expected IOPS could also be that the user is not driving enough I/O to the EBS volumes. The size of the volume may not affect IOPS.

NEW QUESTION 273
A storage admin wants to encrypt all the objects stored in S3 using server side encryption. The user does not want to use the AES 256 encryption key provided by S3. How can the user achieve this?

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

Answer: D

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

NEW QUESTION 277
A user is trying to create a PIOPS EBS volume with 8 GB size and 200 IOPS. Will AWS create the volume?

A. Yes, since the ratio between EBS and IOPS is less than 30
B. No, since the PIOPS and EBS size ratio is less than 30
C. No, the EBS size is less than 10 GB
D. Yes, since PIOPS is higher than 100

Answer: C

Explanation:
A provisioned IOPS EBS volume can range in size from 10 GB to 1 TB and the user can provision up to 4000 IOPS per volume. The ratio of IOPS provisioned to the volume size requested should be a maximum of 30; for example, a volume with 3000 IOPS must be at least 100 GB.

NEW QUESTION 278
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 282
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 as --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 287
A user is trying to create an EBS volume with the highest PIOPS supported by EBS. What is the minimum size of EBS required to have the maximum IOPS?

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

Answer: C

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

NEW QUESTION 288
An organization is trying to create various IAM users. Which of the below mentioned options is not a valid IAM username?

A. John.cloud
B. john@cloud
C. John=cloud
D. john#cloud

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. Whenever the organization is creating an IAM user, there should be a unique ID for each user. The names of users, groups, roles, instance profiles must be alphanumeric, including the following common characters: plus (+., equal (=., comma (,., period (.., at (@., and dash (-..

NEW QUESTION 293
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 errorat 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 297
Which of the below mentioned AWS RDS logs cannot be viewed from the console for MySQL?

A. Error Log
B. Slow Query Log
C. Transaction Log
D. General Log

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 querylog, and general logs. RDS does not support viewing the transaction logs.

NEW QUESTION 299
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 300
A user has created a VPC with public and private subnets using the VPC wizard. Which of the below mentioned statements is true in this scenario?

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

Answer: B

Explanation:
A Virtual Private Cloud (VPC. is a virtual network dedicated to the user??s AWS account. A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance of a smaller or higher size, respectively. The VPC has an implied router and the VPC wizard updates the main route table used with the private subnet, creates a custom route table and associates it with the public subnet.

NEW QUESTION 303
A user has two EC2 instances running in two separate regions. The user is running an internal memory management tool, which captures the data and sends it to CloudWatch in US East, using a CLI with the same namespace and metric. Which of the below mentioned options is true with respect to the above statement?

A. The setup will not work as CloudWatch cannot receive data across regions
B. CloudWatch will receive and aggregate the data based on the namespace and metric
C. CloudWatch will give an error since the data will conflict due to two sources
D. CloudWatch will take the data of the server, which sends the data first

Answer: B

Explanation:
Amazon CloudWatch does not differentiate the source of a metric when receiving custom data. If the user is publishing a metric with the same namespace and dimensions from different sources, CloudWatch will treat them as a single metric. If the data is coming with the same timezone within a minute, CloudWatch will aggregate the data. It treats these as a single metric, allowing the user to get the statistics, such as minimum, maximum, average, and the sum of all across all servers.

NEW QUESTION 308
An organization has created a Queue named ??modularqueue?? with SQS. The organization is not performing any operations such as SendMessage, ReceiveMessage, DeleteMessage, GetQueueAttributes, SetQueueAttributes, AddPermission, and RemovePermission on the queue. What can happen in this scenario?

A. AWS SQS sends notification after 15 days for inactivity on queue
B. AWS SQS can delete queue after 30 days without notification
C. AWS SQS marks queue inactive after 30 days
D. AWS SQS notifies the user after 2 weeks and deletes the queue after 3 weeks.

Answer: B

Explanation:
Amazon SQS can delete a queue without notification if one of the following actions hasn't been performed on it for 30 consecutive days: SendMessage, ReceiveMessage, DeleteMessage, GetQueueAttributes, SetQueueAttributes, AddPermission, and RemovePermission.

NEW QUESTION 309
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 312
An organization is planning to create a user with IAM. They are trying to understand the limitations of IAM so that they can plan accordingly. Which of the below mentioned statements is not true with respect to the limitations of IAM?

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

Answer: A

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The default maximums for each of the IAM entities is given below:
Groups per AWS account: 100 Users per AWS account: 5000 Roles per AWS account: 250
Number of groups per user: 10 (that is, one user can be part of these many groups.

NEW QUESTION 317
A user has created a VPC with two subnets: one public and one private. The user is planning to run the patch update for the instances in the private subnet. How can the instances in the private subnet connect to theinternet?

A. Use the internet gateway with a private IP
B. Allow outbound traffic in the security group for port 80 to allow internet updates
C. The private subnet can never connect to the internet
D. Use NAT with an elastic IP

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 two subnets (one private and one public., he would need a Network Address Translation (NAT. instance with the elastic IP address. This enables the instances in the private subnet to send requests to the internet (for example, to perform software updates..

NEW QUESTION 318
A user is running a batch process on EBS backed EC2 instances. The batch process starts a few instances to process hadoop Map reduce jobs which can run between 50 ?V 600 minutes or sometimes for more time. The user wants to configure that the instance gets terminated only when the process is completed. How can the user configure this with CloudWatch?

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

Answer: D

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

NEW QUESTION 319
A user has created a VPC with public and private subnets using the VPC Wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24. Which of the below mentioned entries are required in the main route table to allow the instances in VPC to communicate with each other?

A. Destination : 20.0.0.0/24 and Target : VPC
B. Destination : 20.0.0.0/16 and Target : ALL
C. Destination : 20.0.0.0/0 and Target : ALL
D. Destination : 20.0.0.0/24 and Target : Local

Answer: D

NEW QUESTION 322
A sysadmin has created the below mentioned policy on an S3 bucket named cloudacademy. The bucket has both AWS.jpg and index.html objects. What does this policy define?
"Statement": [{
"Sid": "Stmt1388811069831",
"Effect": "Allow", "Principal": { "AWS": "*"},
"Action": ["s3:GetObjectAcl", "s3:ListBucket", "s3:GetObject"], "Resource": ["arn:aws:s3:::cloudacademy/*.jpg]
}]

A. It will make all the objects as well as the bucket public
B. It will throw an error for the wrong action and does not allow to save the policy
C. It will make the AWS.jpg object as public
D. It will make the AWS.jpg as well as the cloudacademy bucket as public

Answer: B

NEW QUESTION 323
A user has launched an EC2 instance and deployed a production application in it. The user wants to prohibit any mistakes from the production team to avoid accidental termination. How can the user achieve this?

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

Answer: A

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

NEW QUESTION 324
A user has created a launch configuration for Auto Scaling where CloudWatch detailed monitoring is disabled. The user wants to now enable detailed monitoring. How can the user achieve this?

A. Update the Launch config with CLI to set InstanceMonitoringDisabled = false
B. The user should change the Auto Scaling group from the AWS console to enable detailed monitoring
C. Update the Launch config with CLI to set InstanceMonitoring.Enabled = true
D. Create a new Launch Config with detail monitoring enabled and update the Auto Scaling group

Answer: D

Explanation:
CloudWatch is used to monitor AWS as well as the custom services. To enable detailed instance monitoring for a new Auto Scaling group, the user does not need to take any extra steps. When the user creates the AutoScaling launch config as the first step for creating an Auto Scaling group, each launch configuration contains a flag named InstanceMonitoring.Enabled. The default value of this flag is true. When the user has created a launch configuration with InstanceMonitoring.Enabled = false it will involve multiple steps to enable detail monitoring. The steps are:
Create a new Launch config with detailed monitoring enabled Update the Auto Scaling group with a new launch config Enable detail monitoring on each EC2 instance

NEW QUESTION 325
A user has launched an EC2 instance from an instance store backed AMI. The user has attached an additional instance store volume to the instance. The user wants to create an AMI from the running instance. Will the AMI have the additional instance store volume data?

A. Yes, the block device mapping will have information about the additional instance store volume
B. No, since the instance store backed AMI can have only the root volume bundled
C. It is not possible to attach an additional instance store volume to the existing instance store backed AMI instance
D. No, since this is ephermal storage it will not be a part of the AMI

Answer: A

Explanation:
When the user has launched an EC2 instance from an instance store backed AMI and added an instance store volume to the instance in addition to the root device volume, the block device mapping for the new AMI contains the information for these volumes as well. In addition, the block device mappings for the instances those are launched from the new AMI will automatically contain information for these volumes.

NEW QUESTION 327
......

CertLeader
Leader of IT Certifications

100% Valid and Newest Version SOA-C01 Questions & Answers shared by Certleader
https://www.certleader.com/SOA-C01-dumps.html (639 Q&As)

Thank You for Trying Our Product

* 100% Pass or Money Back

    All our products come with a 90-day Money Back Guarantee.

* One year free update

    You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

    We currently serve more than 30,000,000 customers.

* Shop Securely

    All transactions are protected by VeriSign!

100% Pass Your SOA-C01 Exam with Our Prep Materials Via below:

https://www.certleader.com/SOA-C01-dumps.html

The Leader of IT Certification   *visit - https://www.certleader.com*