Amazon-Web-Services

Exam Questions SOA-C01

AWS Certified SysOps Administrator - Associate

NEW QUESTION 1
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 2
You have a web application leveraging an Elastic Load Balancer (ELB) In front of the web servers deployed using an Auto Scaling Group. Your database is running on Relational Database Service (RDS) The application serves out technical articles and responses to them in general there are more views of an article than there are responses to the article. On occasion, an article on the site becomes extremely popular resulting in significant traffic Increases that causes the site to go down.
What could you do to help alleviate the pressure on the infrastructure while maintaining availability during these events?
Choose 3 answers

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

Answer: ABC

Explanation:
The questions mention RDS so an answer that includes that as part of the solution makes sense. Also, Route53 does nothing to alleviate pressure on the infrastructure, it??s for failover. E is counterproductive. It talks about failing over to an error page on S3.

NEW QUESTION 3
When assessing an organization s use of AWS API access credentials which of the following three credentials should be evaluated? Choose 3 answers

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

Answer: ACD

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

NEW QUESTION 4
Which services allow the customer to retain full administrative privileges of the underlying EC2 instances?
Choose 2 answers

A. Amazon Elastic Map Reduce
B. Elastic Load Balancing
C. AWS Elastic Beanstalk
D. Amazon ElastiCache
E. Amazon Relational Database service

Answer: AC

Explanation:
The below services provide Root level access:
* EC2
* Elastic Beanstalk
* Elastic MapReduce ?V Master Node
* Opswork

NEW QUESTION 5
You have a web-style application with a stateless but CPU and memory-intensive web tier running on
a cc2 8xlarge EC2 instance inside of a VPC The instance when under load is having problems returning requests within the SLA as defined by your business The application maintains its state in a DynamoDB table, but the data tier is properly provisioned and responses are consistently fast.
How can you best resolve the issue of the application responses not meeting your SLA?

A. Add another cc2 8xlarge application instance, and put both behind an Elastic Load Balancer
B. Move the cc2 8xlarge to the same Availability Zone as the DynamoDB table
C. Cache the database responses in ElastiCache for more rapid access
D. Move the database from DynamoDB to RDS MySQL in scale-out read-replica configuration

Answer: C

Explanation:
But it is possibly A as DynamoDB is automatically available across three facilities in an AWS Region. So moving in to a same AZ is not possible / necessary.
In this case the DB layer is not the issue, the EC2 8xlarge is the issue; so add another one with a ELB in-frond of it.
See also: https://aws.amazon.com/dynamodb/faqs/

NEW QUESTION 6
You run a web application where web servers on EC2 Instances are In an Auto Scaling group Monitoring over the last 6 months shows that 6 web servers are necessary to handle the minimum load During the day up to 12 servers are needed Five to six days per year, the number of web servers required might go up to 15.
What would you recommend to minimize costs while being able to provide hill availability?

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

Answer: A

Explanation:
The only plausible answer is A because all other answers include Spot Instances that can be removed without warning and that would not be highly available.

NEW QUESTION 7
You have been asked to propose a multi-region deployment of a web-facing application where a controlled portion of your traffic is being processed by an alternate region.
Which configuration would achieve that goal?

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

Answer: A

Explanation:
The question is asking ??a controlled portion of your traffic??, that would be established with weighted routing policy.
See: http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

NEW QUESTION 8
You are running a database on an EC2 instance, with the data stored on Elastic Block Store (EBS) for persistence. At times throughout the day, you are seeing large variance in the response times of the database queries Looking into the instance with the isolate command you see a lot of wait time on the disk volume that the database's data is stored on.
What two ways can you improve the performance of the database's storage while maintaining the current persistence of the data?
Choose 2 answers

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

Answer: BC

NEW QUESTION 9
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 10
You have decided to change the Instance type for instances running In your application tier that are using Auto Scaling.
In which area below would you change the instance type definition?

A. Auto Scaling launch configuration
B. Auto Scaling group

C. Auto Scaling policy
D. Auto Scaling tags

Answer: A

Explanation:
Reference:
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/WhatIsAutoScaling.html

NEW QUESTION 10
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 14
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 17
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 22
You are tasked with setting up a cluster of EC2 Instances for a NoSQL database. The database requires random read IO disk performance up to a 100,000 IOPS at 4KB block side per node.
Which of the following EC2 instances will perform the best for this workload?

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

Answer: C

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

NEW QUESTION 27
When an EC2 EBS-backed (EBS root) instance is stopped, what happens to the data on any ephemeral store volumes?

A. Data will be deleted and win no longer be accessible
B. Data is automatically saved in an EBS volume.
C. Data is automatically saved as an EBS snapshot
D. Data is unavailable until the instance is restarted

Answer: A

Explanation:
See: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html#instance-store- lifetime
However, data in the instance store is lost under the following circumstances:
?V The underlying disk drive fails
?V The instance stops
?V The instance terminates

NEW QUESTION 31
Your team is excited about the use of AWS because now they have access to ??programmable Infrastructure" You have been asked to manage your AWS infrastructure In a manner similar to the way you might manage application code You want to be able to deploy exact copies of different versions of your infrastructure, stage changes into different environments, revert back to previous versions, and identify what versions are running at any particular time (development, test, QA, production).
Which approach addresses this requirement?

A. Use cost allocation reports and AWS OpsWorks to deploy and manage your infrastructure.
B. Use AWS CloudWatch metrics and alerts along with resource tagging to deploy and manage your infrastructure.
C. Use AWS Beanstalk and a version control system like GIT to deploy and manage your infrastructure.
D. Use AWS CloudFormation and a version control system like GIT to deploy and manage your infrastructure.

Answer: D

Explanation:
Reference:
?V Answer A: does not provide versioning
?V Answer B: does not provide versioning
?V Answer C: Beanstalk provide version control over your application (not infrastructure)
Extract from what is AWS CloudFormation: (http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html)
Easily Control and Track Changes to Your Infrastructure In some cases, you might have underlying resources that you want to upgrade incrementally. For example, you might change to a higher performing instance type in your Auto Scaling launch configuration so that you can reduce the maximum number of instances in your Auto Scaling group. If problems occur after you complete the update, you might need to roll back your infrastructure to the original settings. To do this manually, you not only have to remember which resources were changed, you also have to know what the original settings were.
When you provision your infrastructure with AWS CloudFormation, the AWS CloudFormation template describes exactly what resources are provisioned and their settings. Because these templates are text files, you simply track differences in your templates to track changes to your infrastructure, similar to the way developers control revisions to source code. For example, you can use a version control system with your templates so that you know exactly what changes were made, who made them, and when. If at any point you need to reverse changes to your infrastructure, you can use a previous version of your template.

NEW QUESTION 36
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 38
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 43
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 45
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 49
A user is planning to use AWS Cloud formation for his automatic deployment requirements. Which of the below mentioned components are required as a part of the template?

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

Answer: D

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

NEW QUESTION 54
A user has recently started using EC2. The user launched one EC2 instance in the default subnet in EC2-VPC Which of the below mentioned options is not attached or available with the EC2 instance when it is launched?

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

Answer: C

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

NEW QUESTION 59
A user has launched an EC2 instance. The user is planning to setup the CloudWatch alarm. Which of the below mentioned actions is not supported by the CloudWatch alarm?

A. Notify the Auto Scaling launch config to scale up
B. Send an SMS using SNS
C. Notify the Auto Scaling group to scale down
D. Stop the EC2 instance

Answer: A

Explanation:
A user can create a CloudWatch alarm that takes various actions when the alarm changes state. An alarm watches a single metric over the time period that the user has specified, 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 actions could be sending a notification to an Amazon Simple Notification Service topic (SMS, Email, and HTTP end point notifying the Auto Scaling policy or changing the state of the instance to Stop/Terminate.
CloudWatch cannot change the auto-scaling launch configuration.
B ?V It can send an SMS with SNS
C ?V Auto-scaling uses CloudWatch metrics to scale up and down.
D ?V CloudWatch can stop instances

NEW QUESTION 60
A sysadmin has created the below mentioned policy and applied to an S3 object named aws.jpg. The aws.jpg is inside a bucket named cloudacademy. What does this policy define?
"Statement": [{
"Sid": "Stmt1388811069831",
"Effect": "Allow", "Principal": { "AWS": "*"},
"Action": ["s3:GetObjectAcl", "s3:ListBucket", "s3:GetObject"], "Resource": ["arn:aws:s3:::cloudacademy/*.jpg"]
}]

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

Answer: A

NEW QUESTION 63
A user is trying to save some cost on the AWS services. Which of the below mentioned options will not help him save cost?

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

Answer: B

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

NEW QUESTION 65
A user has created a subnet with VPC and launched an EC2 instance in that subnet with only default settings. Which of the below mentioned options is ready to use on the EC2 instance as soon as it is launched?

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

Answer: B

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

NEW QUESTION 70
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 72
A user has configured the Auto Scaling group with the minimum capacity as 3 and the maximum capacity as 5. When the user configures the AS group, how many instances will Auto Scaling launch?

A. 3
B. 5
C. 2

Answer: A

Explanation:

The default launch is going to configure 3 instances. It will only go to 5 under certain conditions. It won??t launch initially 5 instances. This question doesn??t give enough information to answer fully.

NEW QUESTION 74
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 76
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 78
A user has created an ELB with three instances. How many security groups will ELB create by default?

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

Answer: C

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

NEW QUESTION 81
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 82
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 86
A root AWS account owner is trying to understand various options to set the permission to AWS S3. Which of the below mentioned options is not the right option to grant permission for S3?

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

Answer: B

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

NEW QUESTION 90
A sysadmin has created a shopping cart application and hosted it on EC2. The EC2 instances are running behind ELB. The admin wants to ensure that the end user request will always go to the EC2 instance where the user session has been created. How can the admin configure this?

A. Enable ELB cross zone load balancing
B. Enable ELB cookie setup
C. Enable ELB sticky session
D. Enable ELB connection draining

Answer: C

Explanation:
Generally, AWS ELB routes each request to a zone with the minimum load. The Elastic Load Balancer provides a feature called sticky session which binds the user??s session with a specific EC2 instance. If the sticky session is enabled the first request from the user will be redirected to any of the EC2 instances. But, henceforth, all requests from the same user will be redirected to the same EC2 instance. This ensures that all requests coming from the user during the session will be sent to the same application instance.

NEW QUESTION 95
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 98
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 103
A user has configured the AWS CloudWatch alarm for estimated usage charges in the US East region. Which of the below mentioned statements is not true with respect to the estimated charges?
Exhibit:

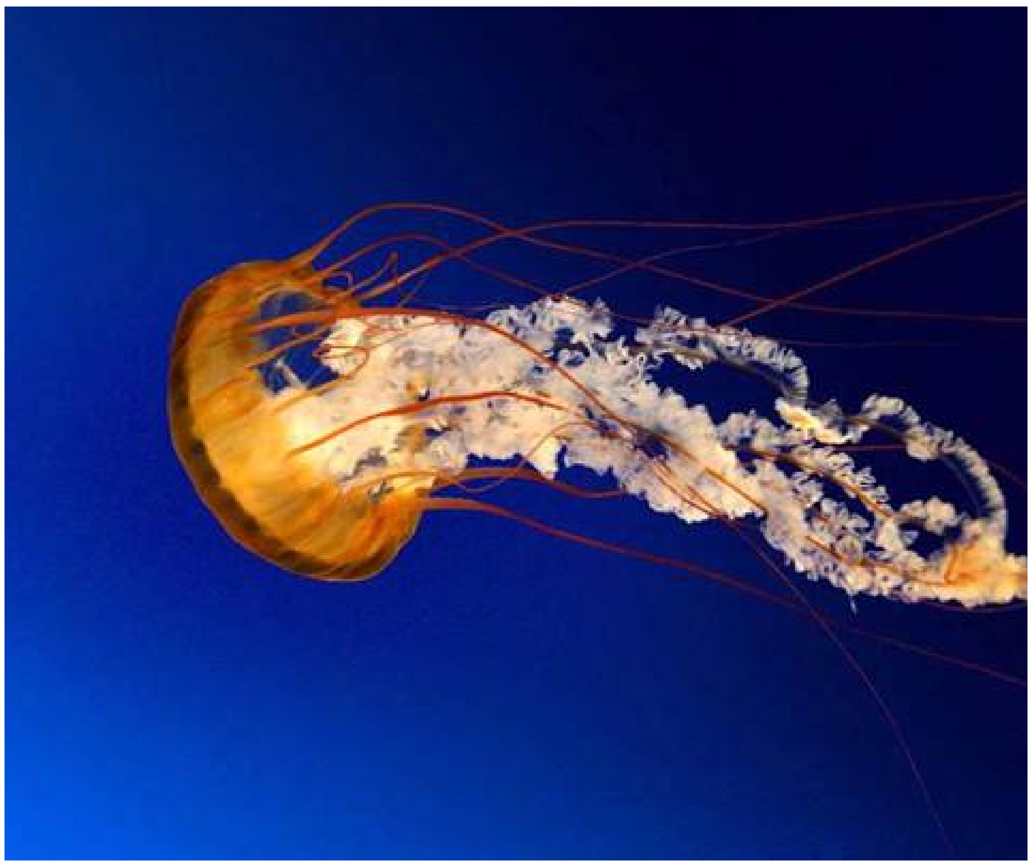

A. It will store the estimated charges data of the last 14 days
B. It will include the estimated charges of every AWS service
C. The metric data will represent the data of all the regions
D. The metric data will show data specific to that region

Answer: D

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

NEW QUESTION 107
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 109
A user has enabled the Multi AZ feature with the MS SQL RDS database server. Which of the below mentioned statements will help the user understand the Multi AZ feature better?

A. In a Multi AZ, AWS runs two DBs in parallel and copies the data asynchronously to the replica copy
B. In a Multi AZ, AWS runs two DBs in parallel and copies the data synchronously to the replica copy
C. In a Multi AZ, AWS runs just one DB but copies the data synchronously to the standby replica
D. AWS MS SQL does not support the Multi AZ feature

Answer: C

Explanation:
Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy, eliminate I/O freezes, and minimize latency spikes during system backups. Running a DB instance with high availability can enhance availability during planned system maintenance, and help protect your databases against DB instance failure and Availability Zone disruption.Note that the high-availability feature is not a scaling solution for read-only scenarios; you cannot use a standby replica to serve read traffic. To service

read-only traffic, you should use a read replica.

NEW QUESTION 113
A user is publishing custom metrics to CloudWatch. Which of the below mentioned statements will help the user understand the functionality better?

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

Answer: B

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

NEW QUESTION 118
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 121
A user has setup Auto Scaling with ELB on the EC2 instances. The user wants to configure that
whenever the CPU utilization is below 10%, Auto Scaling should remove one instance. How can the user configure this?

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

Answer: D

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

NEW QUESTION 122
A user has enabled detailed CloudWatch metric monitoring on an Auto Scaling group. Which of the below mentioned metrics will help the user identify the total number of instances in an Auto Scaling group cluding pending, terminating and running instances?

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

Answer: A

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

NEW QUESTION 125
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 130
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 134
A root account owner has created an S3 bucket testmycloud. The account owner wants to allow everyone to upload the objects as well as enforce that the person who uploaded the object should manage the permission of those objects. Which is the easiest way to achieve this?

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

Answer: C

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

NEW QUESTION 138
An organization has setup consolidated billing with 3 different AWS accounts. Which of the below mentioned advantages will organization receive in terms of the AWS pricing?

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

Answer: B

Explanation:
AWS consolidated billing enables the organization to consolidate payments for multiple Amazon Web Services (AWS. accounts within a single organization by making a single paying account. For billing purposes, AWS treats all the accounts on the consolidated bill as one account. Some services, such as Amazon EC2 and Amazon S3 have volume pricing tiers across certain usage dimensions that give the user lower prices when he uses the service more.

NEW QUESTION 140
A user has setup an RDS DB with Oracle. The user wants to get notifications when someone modifies the security group of that DB. How can the user configure that?

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

Answer: C

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

NEW QUESTION 141

A user is trying to setup a recurring Auto Scaling process. The user has setup one process to scale up every day at 8 am and scale down at 7 PM. The user is trying to setup another recurring process which scales up on the 1st of every month at 8 AM and scales down the same day at 7 PM. What will Auto Scaling do in this scenario?

A. Auto Scaling will execute both processes but will add just one instance on the 1st
B. Auto Scaling will add two instances on the 1st of the month
C. Auto Scaling will schedule both the processes but execute only one process randomly
D. Auto Scaling will throw an error since there is a conflict in the schedule of two separate Auto Scaling Processes

Answer: D

Explanation:
Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. The user can also configure the recurring schedule action which will follow the Linux cron format. As per Auto Scaling, a scheduled action must have a unique time value. If the user attempts to schedule an activity at a time when another existing activity is already scheduled, the call will be rejected with an error message noting the conflict.

NEW QUESTION 145
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 146
You are managing the AWS account of a big organization. The organization has more than 1000+ employees and they want to provide access to the various services to most of the employees. Which of the below mentioned options is the best possible solution in this case?

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

Answer: D

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

NEW QUESTION 149
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 151
A user has created a Cloudformation stack. The stack creates AWS services, such as EC2 instances, ELB, AutoScaling, and RDS. While creating the stack it created EC2, ELB and AutoScaling but failed to
create RDS. What will Cloudformation do in this scenario?

A. Cloudformation can never throw an error after launching a few services since it verifies all the steps before launching
B. It will warn the user about the error and ask the user to manually create RDS
C. Rollback all the changes and terminate all the created services
D. It will wait for the user??s input about the error and correct the mistake after the input

Answer: C

Explanation:

AWS Cloudformation is an application management tool which provides application modelling, deployment, configuration, management and related activities. The AWS Cloudformation stack is a collection of AWS resources which are created and managed as a single unit when AWS CloudFormation instantiates a template. If any of the services fails to launch, Cloudformation will rollback all the changes and terminate or delete all the created services.

NEW QUESTION 154
A user has created a VPC with public and private subnets using the VPC wizard. The user has not launched any instance manually and is trying to delete the VPC. What will happen in this scenario?

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

Answer: D

Explanation:
A Virtual Private Cloud (VPC. is a virtual network dedicated to the user??s AWS account. A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance with an elastic IP. If the user is trying to delete the VPC it will not allow as the NAT instance is still running.

NEW QUESTION 155
A user has launched an EBS backed instance with EC2-Classic. The user stops and starts the instance. Which of the below mentioned statements is not true with respect to the stop/start action?

A. The instance gets new private and public IP addresses
B. The volume is preserved
C. The Elastic IP remains associated with the instance
D. The instance may run on a anew host computer

Answer: C

Explanation:
A user can always stop/start an EBS backed EC2 instance. When the user stops the instance, it first enters the stopping state, and then the stopped state. AWS does not charge the running cost but charges only for the EBS storage cost. If the instance is running in EC2-Classic, it receives a new private IP address; as the Elastic IP address (EIP. associated with the instance is no longer associated with that instance.

NEW QUESTION 158
A user has launched an RDS postgreSQL DB with AWS. The user did not specify the maintenance window during creation. The user has configured RDS to update the DB instance type from micro to large. If the user wants to have it during the maintenance window, what will AWS do?

A. AWS will not allow to update the DB until the maintenance window is configured
B. AWS will select the default maintenance window if the user has not provided it
C. AWS will ask the user to specify the maintenance window during the update
D. It is not possible to change the DB size from micro to large with RDS

Answer: B

Explanation:
AWS RDS has a compulsory maintenance window which by default is 30 minutes. If the user does not specify the maintenance window during the creation of RDS then AWS will select a 30-minute maintenance window randomly from an 8-hour block of time per region. In this case, Amazon RDS assigns a 30-minute maintenance window on a randomly selected day of the week.

NEW QUESTION 160
A user has created a VPC with CIDR 20.0.0.0/16 using VPC Wizard. The user has created a public CIDR (20.0.0.0/24. and a VPN only subnet CIDR (20.0.1.0/24. along with the hardware VPN access to connect to the user??s data centre. Which of the below mentioned components is not present when the VPC is setup with the wizard?

A. Main route table attached with a VPN only subnet
B. A NAT instance configured to allow the VPN subnet instances to connect with the internet
C. Custom route table attached with a public subnet
D. An internet gateway for a public subnet

Answer: B

Explanation:
The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will update the main route table used with the VPN-only subnet, create a custom route table and associate it with the public subnet. It also creates an internet gateway for the public subnet. The wizard does not create a NAT instance by default. The user can create it manually and attach it with a VPN only subnet.

NEW QUESTION 165
A user has setup an Auto Scaling group. The group has failed to launch a single instance for more than 24 hours. What will happen to Auto Scaling in this condition?

A. Auto Scaling will keep trying to launch the instance for 72 hours
B. Auto Scaling will suspend the scaling process
C. Auto Scaling will start an instance in a separate region

D. The Auto Scaling group will be terminated automatically

Answer: B

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

NEW QUESTION 166
A sysadmin has created the below mentioned policy on an S3 bucket named cloudacademy. What does this policy define?
"Statement": [{
"Sid": "Stmt1388811069831",
"Effect": "Allow", "Principal": { "AWS": "*"},
"Action": ["s3:GetObjectAcl", "s3:ListBucket"], "Resource": ["arn:aws:s3:::cloudacademy]
}]

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

Answer: D

Explanation:
A sysadmin can grant permission to the S3 objects or the buckets to any user or make objects public using the bucket policy and user policy. Both use the JSON-based access policy language. Generally if the user is defining the ACL on the bucket, the objects in the bucket do not inherit it and vice 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. In the sample policy the action says ??S3:ListBucket?? for effect Allow on Resource arn:aws:s3:::cloudacademy. This will make the cloudacademy bucket public.
"Statement": [{
"Sid": "Stmt1388811069831",
"Effect": "Allow", "Principal": { "AWS": "*" },
"Action": ["s3:GetObjectAcl", "s3:ListBucket"], "Resource": ["arn:aws:s3:::cloudacademy]
}]

NEW QUESTION 167
A user is trying to connect to a running EC2 instance using SSH. However, the user gets a Host key not found error. Which of the below mentioned options is a possible reason for rejection?

A. The user has provided the wrong user name for the OS login
B. The instance CPU is heavily loaded
C. The security group is not configured properly
D. The access key to connect to the instance is wrong

Answer: A

Explanation:
If the user is trying to connect to a Linux EC2 instance and receives the Host Key not found error the probable reasons are:
The private key pair is not right The user name to login is wrong

NEW QUESTION 169
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 172
A user has created a queue named ??awsmodule?? with SQS. One of the consumers of queue is down for 3 days and then becomes available. Will that component receive message from queue?

A. Yes, since SQS by default stores message for 4 days
B. No, since SQS by default stores message for 1 day only
C. No, since SQS sends message to consumers who are available that time
D. Yes, since SQS will not delete message until it is delivered to all consumers

Answer: A

Explanation:
SQS allows the user to move data between distributed components of applications so they can perform different tasks without losing messages or requiring each component to be always available. Queues retain messages for a set period of time. By default, a queue retains messages for four days. However, the user can configure a queue to retain messages for up to 14 days after the message has been sent.

NEW QUESTION 176
An organization has setup multiple IAM users. The organization wants that each IAM user accesses the IAM console only within the organization and not from outside. How can it achieve this?

A. Create an IAM policy with the security group and use that security group for AWS console login
B. Create an IAM policy with a condition which denies access when the IP address range is not from the organization
C. Configure the EC2 instance security group which allows traffic only from the organization??s IP range
D. Create an IAM policy with VPC and allow a secure gateway between the organization and AWS Console

Answer: B

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 many other parameters. If the organization wants the user to access only from a specific IP range, they should set an IAM policy condition which denies access when the IP is not in a certain range. E.g. The sample policy given below denies all traffic when the IP is not in a certain range.

```
"Statement": [{
"Effect": "Deny",
"Action": "*",
"Resource": "*", "Condition": { "NotIpAddress": {
"aws:SourceIp": ["10.10.10.0/24", "20.20.30.0/24"]
}
}
}]
```

NEW QUESTION 181
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 183

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

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

Answer: A

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

NEW QUESTION 184
A user has configured ELB with a TCP listener at ELB as well as on the back-end instances. The user wants to enable a proxy protocol to capture the source and destination IP information in the header. Which of the below mentioned statements helps the user understand a proxy protocol with TCP configuration?

A. If the end user is requesting behind a proxy server then the user should not enable a proxy protocol on ELB
B. ELB does not support a proxy protocol when it is listening on both the load balancer and the back- end instances
C. Whether the end user is requesting from a proxy server or directly, it does not make a difference for the proxy protocol
D. If the end user is requesting behind the proxy then the user should add the ??isproxy?? flag to the ELB Configuration

Answer: A

Explanation:
When the user has configured Transmission Control Protocol (TCP. or Secure Sockets Layer (SSL. for both front-end and back-end connections of the Elastic Load Balancer, the load balancer forwards the request to the back-end instances without modifying the request headers unless the proxy header is enabled. If the end user is requesting from a Proxy Protocol enabled proxy server, then the ELB admin should not enable the Proxy Protocol on the load balancer. If the Proxy Protocol is enabled on both the proxy server and the load balancer, the load balancer will add another header to the request which already has a header from the proxy server. This duplication may result in errors.

NEW QUESTION 189
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 190
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 195
A user has enabled versioning on an S3 bucket. The user is using server side encryption for data at Rest. If the user is supplying his own keys for encryption (SSE-C., which of the below mentioned statements is true?

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

Answer: B

Explanation:
AWS S3 supports client side or server side encryption to encrypt all data at rest. The server side encryption can either have the S3 supplied AES-256 encryption

key or the user can send the key along with each API call to supply his own encryption key (SSE-C.. If the bucket is versioning-enabled, each object version uploaded by the user using the SSE-C feature can have its own encryption key. The user is responsible for tracking which encryption key was used for which object's version

NEW QUESTION 200
A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24. What will happen in this scenario?

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

Answer: D

Explanation:
A Virtual Private Cloud (VPC. is a virtual network dedicated to the user??s AWS account. A user can create a subnet with VPC and launch instances inside that subnet. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet.

NEW QUESTION 201
A root account owner is trying to understand the S3 bucket ACL. Which of the below mentioned options cannot be used to grant ACL on the object using the authorized predefined group?

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

Answer: D

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

NEW QUESTION 205
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 208
A user has setup a VPC with CIDR 20.0.0.0/16. The VPC has a private subnet (20.0.1.0/24. and a public subnet (20.0.0.0/24.. The user??s data centre has CIDR of 20.0.54.0/24 and 20.1.0.0/24. If the private subnet wants to communicate with the data centre, what will happen?

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

Answer: D

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

NEW QUESTION 211
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 212
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 217
A user has created a VPC with the public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The public subnet uses CIDR 20.0.1.0/24. The user is planning to host a web server in the public subnet (port 80. and a DB server in the private subnet (port 3306.. The user is configuring a security group for the public subnet (WebSecGrp. and the private subnet (DBSecGrp.. Which of the below mentioned entries is required in the web server security group (WebSecGrp.?

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

Answer: A

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

NEW QUESTION 222
A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services provides detailed monitoring with CloudWatch without charging the user extra?

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

Answer: B

Explanation:
CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. Services, such as RDS, ELB, OpsWorks, and Route 53 can provide the monitoring data every minute without charging the user.

NEW QUESTION 223
A user has configured Auto Scaling with 3 instances. The user had created a new AMI after updating one of the instances. If the user wants to terminate two specific instances to ensure that Auto Scaling launches an instances with the new launch configuration, which command should he run?

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

Answer: D

Explanation:
The Auto Scaling command as-terminate-instance-in-auto-scaling-group <Instance ID> will terminate the specific instance ID. The user is required to specify the parameter as ?Vno-decrement-desired- capacity to ensure that it launches a new instance from the launch config after terminating the instance. If the user specifies the parameter --decrement-desired-capacity then Auto Scaling will terminate the instance and decrease the desired capacity by 1.

NEW QUESTION 227
A user has created a VPC with CIDR 20.0.0.0/24. The user has used all the IPs of CIDR and wants to increase the size of the VPC. The user has two subnets: public (20.0.0.0/28. and private (20.0.1.0/28.. How can the user change the size of the VPC?

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

Answer: B

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

NEW QUESTION 228
A user has granted read/write permission of his S3 bucket using ACL. Which of the below mentioned options is a valid ID to grant permission to other AWS accounts (grantee. using ACL?

A. IAM User ID
B. S3 Secure ID
C. Access ID
D. Canonical user ID

Answer: D

Explanation:
An S3 bucket ACL grantee can be an AWS account or one of the predefined Amazon S3 groups. The user can grant permission to an AWS account by the email address of that account or by the canonical user ID. If the user provides an email in the grant request, Amazon S3 finds the canonical user ID for that account and adds it to the ACL. The resulting ACL will always contain the canonical user ID for the AWS account, and not the AWS account's email address.

NEW QUESTION 232
A user has launched an EC2 instance store backed instance in the US-East-1a zone. The user created AMI #1 and copied it to the Europe region. After that, the user made a few updates to the application running in the US-East-1a zone. The user makes an AMI#2 after the changes. If the user launches a new instance in Europe from the AMI #1 copy, which of the below mentioned statements is true?

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

Answer: D

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

NEW QUESTION 234
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 239
A sys admin is trying to understand the sticky session algorithm. Please select the correct sequence of steps, both when the cookie is present and when it is not, to help the admin understand the implementation of the sticky session:
ELB inserts the cookie in the response
ELB chooses the instance based on the load balancing algorithm Check the cookie in the service request
The cookie is found in the request
The cookie is not found in the request

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

Answer: C

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

NEW QUESTION 242
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 244
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 246
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 251
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 255
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 259
A user wants to upload a complete folder to AWS S3 using the S3 Management console. How can the user perform this activity?

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

Answer: D

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

NEW QUESTION 264
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 266
The CFO of a company wants to allow one of his employees to view only the AWS usage report page. Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

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

Answer: C

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

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

NEW QUESTION 269
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 273
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 274
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 276
A user has configured an EC2 instance in the US-East-1a zone. The user has enabled detailed monitoring of the instance. The user is trying to get the data from CloudWatch using a CLI. Which of the below mentioned CloudWatch endpoint URLs should the user use?

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

Answer: A

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

NEW QUESTION 281
A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling AddToLoadBalancer
(which adds instances to the load balancer. process for a while. What will happen to the instances launched during the suspension period?

A. The instances will not be registered with ELB and the user has to manually register when the process is resumed
B. The instances will be registered with ELB only once the process has resumed
C. Auto Scaling will not launch the instance during this period due to process suspension
D. It is not possible to suspend only the AddToLoadBalancer process

Answer: A

Explanation:
Auto Scaling performs various processes, such as Launch, Terminate, add to Load Balancer etc. The user can also suspend the individual process. The AddToLoadBalancer process type adds instances to the load balancer when the instances are launched. If this process is suspended, Auto Scaling will launch the instances but will not add them to the load balancer. When the user resumes this process, Auto Scaling will resume adding new instances launched after resumption to the load balancer. However, it will not add running instances that were launched while the process was suspended; those instances must be added manually.

NEW QUESTION 285
A sys admin has enabled a log on ELB. Which of the below mentioned activities are not captured by the log?

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

Answer: B

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

NEW QUESTION 290
A user has moved an object to Glacier using the life cycle rules. The user requests to restore the archive after 6 months. When the restore request is completed the user accesses that archive. Which of the below mentioned statements is not true in this condition?

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

Answer: B

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

NEW QUESTION 292
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 297
A user has enabled versioning on an S3 bucket. The user is using server side encryption for data at rest. If the user is supplying his own keys for encryption (SSE-C., what is recommended to the user for the purpose of security?

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

Answer: D

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

NEW QUESTION 301
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 305
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 309
A user has created an EBS volume of 10 GB and attached it to a running instance. The user is trying to access EBS for first time. Which of the below mentioned options is the correct statement with respect to a first time EBS access?

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

Answer: B

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

NEW QUESTION 312
A user has enabled termination protection on an EC2 instance. The user has also set Instance initiated shutdown behaviour to terminate. When the user shuts down the instance from the OS, what will happen?

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

Answer: B

Explanation:
It is always possible that someone can terminate an EC2 instance using the Amazon EC2 console, command line interface or API by mistake. If the admin wants to prevent the instance from being accidentally terminated, he can enable termination protection for that instance. The user can also setup shutdown behaviour for an EBS backed instance to guide the instance on what should be done when he initiates shutdown from the OS using Instance initiated shutdown behaviour. If the instance initiated behaviour is set to terminate and the user shuts off the OS even though termination protection is enabled, it will still terminate the instance.

NEW QUESTION 317
How can you secure data at rest on an EBS volume?

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

Answer: E

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

NEW QUESTION 322
In order to optimize performance for a compute cluster that requires low inter-node latency, which feature in the following list should you use?

A. AWS Direct Connect
B. Placement Groups
C. VPC private subnets

D. EC2 Dedicated Instances
E. Multiple Availability Zones

Answer: B

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 326
You have a proprietary data store on-premises that must be backed up daily by dumping the data store contents to a single compressed 50GB file and sending the file to AWS. Your SLAs state that any dump file backed up within the past 7 days can be retrieved within 2 hours. Your compliance department has stated that all data must be held indefinitely. The time required to restore the data store from a backup is approximately 1 hour. Your on-premise network connection is capable of sustaining 1gbps to AWS.
Which backup methods to AWS would be most cost-effective while still meeting all of your requirements?

A. Send the daily backup files to Glacier immediately after being generated
B. Transfer the daily backup files to an EBS volume in AWS and take daily snapshots of the volume
C. Transfer the daily backup files to S3 and use appropriate bucket lifecycle policies to send to Glacier
D. Host the backup files on a Storage Gateway with Gateway-Cached Volumes and take daily snapshots

Answer: D

Explanation:
Reference:
http://aws.amazon.com/storagegateway/faqs/

NEW QUESTION 331
You run a web application with the following components Elastic Load Balancer (EL8), 3 Web/Application servers, 1 MySQL RDS database with read replicas, and Amazon Simple Storage Service (Amazon S3) for static content. Average response time for users is increasing slowly.
What three CloudWatch RDS metrics will allow you to identify if the database is the bottleneck? Choose 3 answers

A. The number of outstanding IOs waiting to access the disk.
B. The amount of write latency.
C. The amount of disk space occupied by binary logs on the master.
D. The amount of time a Read Replica DB Instance lags behind the source DB Instance
E. The average number of disk I/O operations per second.

Answer: ABE

NEW QUESTION 334
How can software determine the public and private IP addresses of the Amazon EC2 instance that it is running on?

A. Query the local instance metadata.
B. Query the appropriate Amazon CloudWatch metric.
C. Query the local instance userdata.
D. Use ipconfig or ifconfig command.

Answer: B

NEW QUESTION 335
Which of the following are true regarding encrypted Amazon Elastic Block Store (EBS) volumes? Choose 2 answers

A. Supported on all Amazon EBS volume types
B. Snapshots are automatically encrypted
C. Available to all instance types
D. Existing volumes can be encrypted
E. shared volumes can be encrypted

Answer: AB

Explanation:
This feature is supported on all Amazon EBS volume types (General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic). You can access encrypted Amazon EBS volumes the same way you access existing volumes; encryption and decryption are handled transparently and they require no additional action from you, your Amazon EC2 instance, or your application. Snapshots of encrypted Amazon EBS volumes are automatically encrypted, and volumes that are created from encrypted Amazon EBS snapshots are also automatically encrypted.
Reference: http://docs.aws.amazon.com/kms/latest/developerguide/services-ebs.html

NEW QUESTION 340
Which features can be used to restrict access to data in S3? Choose 2 answers

A. Set an S3 ACL on the bucket or the object.
B. Create a CloudFront distribution for the bucket.
C. Set an S3 bucket policy.
D. Enable IAM Identity Federation
E. Use S3 Virtual Hosting

Answer: AC

Explanation:
https://aws.amazon.com/s3/faqs/

NEW QUESTION 344
The Database Administrator learn is interested in performing manual backups of Amazon DRS Oracle DB instance.
What step be taken to perform the backups?

A. Attach an Amazon EBS volume with Oracle RMAN installed to the RDS instance
B. Take a snapshot of the EBS volume that is attached to the DB instance.
C. Install Oracle Secure Backup on the RDS instance and back up the Oracle database to Amazon S3
D. Take a snapshot of the DB instance

Answer: D

NEW QUESTION 349
A company website hosts patches for software that is sold globally. The website rules in AWS perform will until large software patch is released. The flood of download puts a strain on the web servers and leads to a poor customer experience.
What can the SysOps Administrator propose to enhance customer experience, create a more available platform, and keep costs low?

A. Use an Amazon Cloud Front distribution to cache static content, including software patches.
B. Increase the size of the NAT instance to improve through.
C. Scale out the web servers in advance of patch releases to reduce Auto Scaling delays.
D. Move the content to IO1 and provision additional IOPS to the volume that contains the software patches.

Answer: D

NEW QUESTION 354
A SysOps Administrator has attempted to copy an Marketplace AMI an associated billing Product code that was shared another account. When the copy process is attempted, it fails.
What action can be taken to successfully copy the AMI to the target destination?

A. Use an EC2 instance in the account by using the shared AMI and then created an AMI from the instance
B. Launch an EC2 instance in the account by using the shared AMI and then create an AMI from the instance
C. Use the AWS CLI with the --nobillingProduct flag to execute the copy and ignore the billingProductcode.
D. Create a VPC peering connection between the source and target account to facilitate the AMI copy process.

Answer: D

NEW QUESTION 356
A company three-tier web application is not performing as well as expected. A manager has asked a System Administrator to analyser all the system involved and identity where the performance bottleneck exist.
Which AWS service can be help find bottleneck?

A. Analyse AWS ClouldTrail logs to see which API call are taking the longest to execute
B. Run a performance trace using Amazon Inspector to measure response tone between various API calls
C. Create a rule in AWS Config to send an alert when the performance s noncompliant for each of the tiers
D. Create an Amazon CloudWatch dashboard that contains Amazon EC2 and Amazon RDS metrics

Answer: D

Explanation:
Check the CloudWatch Latency metric
The Latency metric represents the time elapsed, in seconds, after the request leaves the load balancer until a response is received by the load balancer from a registered instance. The preferred statistic for this metric is average, which reports average latency for all requests. A high Latency average value typically indicates a problem with the backend server(s) rather than a problem with the load balancer. Check the maximum statistic to determine the number of latency data points that reach or exceed the load balancer idle timeout value. When latency data points meet or exceed the idle timeout value, it is likely that some requests are timing out, which initiates an HTTP 504 response to clients.

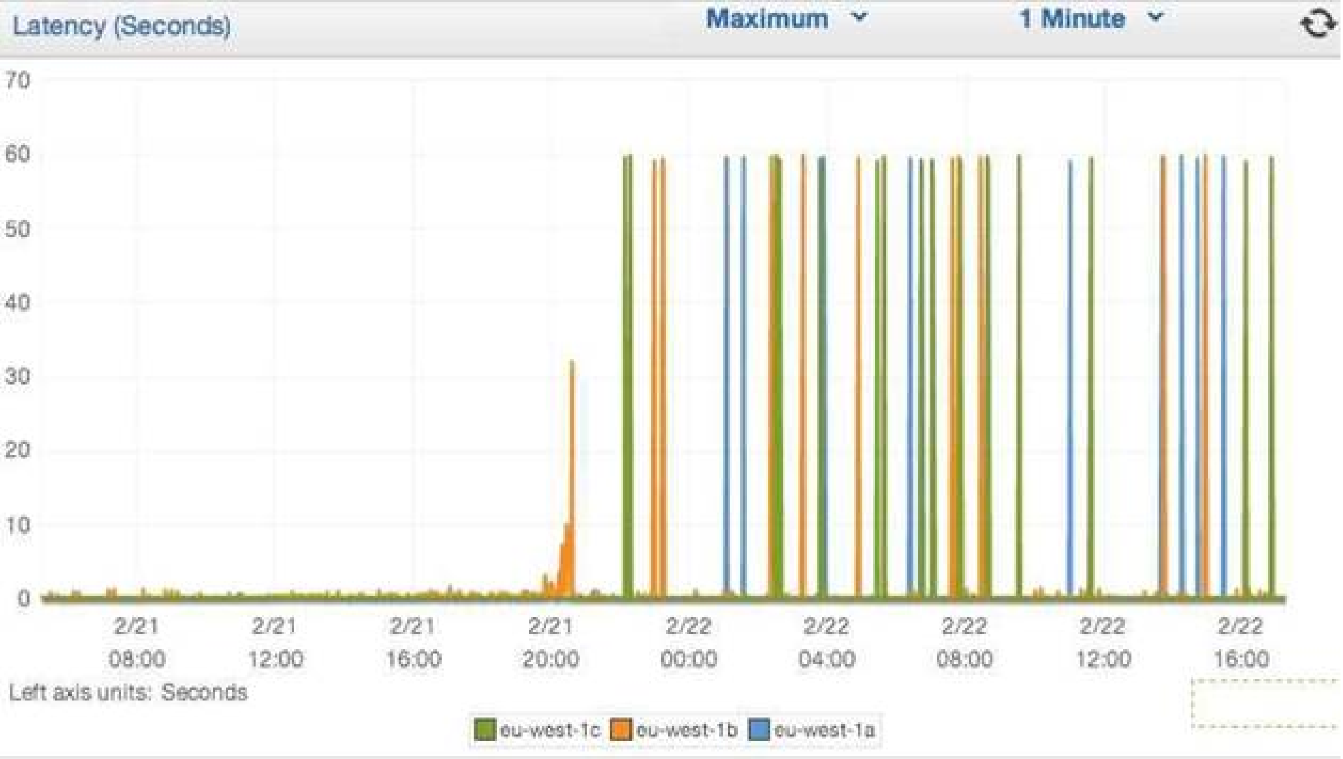

NEW QUESTION 359
A company Development team to access the AWS Management Console. A System Administrator has been asked to find a solution so that the Developers can sign in to the console using Active Directory (AD) credentials and not as IAM users.
What steps should the Systems Administrator take to enable functionality?

A. Set up an Amazon Cognit federation, and the obtain temporary credentials using AWS Security Token Servic
B. Assign the temporary credentials to an IAM role to allow a developers access to the AWS resource.
C. Set up Active Directory Connector to use the corporate AD servers Enable AWS console access under the AWS Directory Service Console for the AD Connector that was just create
D. Created a role with the resources and permissions that the Development team should have access to use.
E. Connect the corporate AD servers to AWS using Amazon Cognito user pools Enable AWS console access within conito, and then assign the appropriate role to the user pool.
F. Create a SAML template file using IAM assign the template to the corporate AD through the Simple AD Grant the Development team access to the SAML template.

Answer: A

NEW QUESTION 360
A SysOps Administrator must take a team's single existing AWS CloudFormation template and split it into smaller, service specific template. All of the service in the template reference a single, shared Amazon S3 bucket.
What should the Administrator do to ensure that this S3 bucket can be referenced by all the service templates?

A. Include the S3 bucket as a mapping in each template
B. Add the S3 bucket as a resource in each template
C. Create the S3 bucket in its own template and export it
D. Generate the S3 bucket using StackSets

Answer: D

NEW QUESTION 364
An organization created an Amazon Elastic File System (Amazon EFS) volume with a file system ID of fs-85baf1fc, and it is actively used by 10 Amazon EC2 hosts.
The organization has become concerned that the file system is not encrypted. How can this be resolved?

A. Enable encryption on each hosts connection to the Amazon EFS volume Each connection must be recreated for encryption to take effect
B. Enable encryption on the existing EFS volume by using the AWS Command Line Interface
C. Enable encryption on each host's local drive Restart each host to encrypt the drive
D. Enable encryption on a newly created volume and copy all data from the original volume Reconnect each host to the new volume

Answer: A

Explanation:
https://docs.aws.amazon.com/efs/latest/ug/encryption.html https://aws.amazon.com/premiumsupport/knowledge-center/encrypt-data-efs/

NEW QUESTION 368
An Amazon EC2 instance is unable to connect to an SMTP server in a different subnet. Other instances are successfully communication with the SMTP servers, however Flow Logs have been enabled on the SMTP server's network interface and show the following information

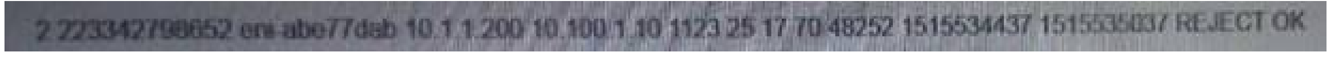

```
2 223342798652 eni-abe77dab 10.1.1.200 10.100.1.10 1123 25 17 70 48252 1515534437 1515535037 REJECT OK
```

A. Add the instance to the security group for the SMTP server and ensure that it is permitted to communicate over TCP port 25.
B. Disable the iptables server on the SMTP server so that the instance can properly communicate over the network.
C. Install an email on the instance to ensure that it communicates correctly on TCP port 25 to theSMTP server.
D. Add a rule to the security group for the instance to explicit permit TCP port 25 outbound to any address.

Answer: D

NEW QUESTION 370
A company has created a separate AWS account for all development work to protect the production environment in this development account, developers have permission to manipulate IAM policies and roles. Corporate policies require that developers and blocked from accessing some services. What is the BEST way to grant the developers privileges in the development account while still complying with corporate policies?

A. Create a service control policy in AWS Organizations and apply it to the development account
B. Create a customer managed policy in IAM and apply it to all users within the development account
C. Create a job function policy in IAM and apply it to all users within the development account
D. Create an IAM policy and apply it in API Gateway to restrict the development account

Answer: A

Explanation:
https://aws.amazon.com/blogs/security/how-to-use-service-control-policies-in-aws-organizations/

NEW QUESTION 375
An errant process is known to use in an entire processor and run at 100%. A SysOps Administrator wants to automate restarting the instance once the problem occurs for more than minutes.
How can this be accomplished?

A. Create an Amazon CloudWatch alarm or the Amazon EC2 instance with basic monitoring Enable an action to restart the instance
B. Create a CloudWatch alarm for the EC2 instance with detailed monitoring Enable an action to restart the instance
C. Create an AWS Lambda function to restart the EC2 instance triggered on a scheduled basis every 2 minutes
D. Create a Lambda function start the EC2 instance triggered by EC2 health

Answer: D

Explanation:
You can use CloudWatch Events to trigger an AWS Lambda function to start and stop your EC2 instances at scheduled intervals.
Note: This article provides an example for a simple solution. For a more robust solution, see AWS Instance Scheduler.
Resolution
CloudWatch Events allows you to create an event that is triggered at a specified time or interval in response to events that take place in your account. For example, you can create an event using CloudWatch Events for a specific time of day, or you can create an alarm when CPU utilization for an instance reaches a specific threshold. You can also configure a Lambda function to start and stop instances when triggered by these events.
In this example, we use Lambda functions to start and stop EC2 instances, and then we use CloudWatch Events to start instances in the morning and stop the instances at night.
1. Open the AWS Lambda console, and choose Create function.
2. Choose Author from scratch.
3. Enter a Name for your function, such as "StopEC2Instances."
4. From the Runtime drop-down menu, choose Python2.7.
5. Expand the Role drop-down menu, and then choose Create a custom role. This opens a new tab or window in your browser.
6. In the IAM Role drop-down menu, choose Create a new IAM Role, and enter a Role Name, such as ??lambda_start_stop_ec2."
7. Expand View Policy Document, choose Edit, and then choose Ok when prompted to read the documentation.

NEW QUESTION 376
......