Exam Questions AWS-SysOps

Amazon AWS Certified SysOps Administrator - Associate

https://www.2passeasy.com/dumps/AWS-SysOps/

NEW QUESTION 1
- (Topic 1)
What would happen to an RDS (Relational Database Service) multi-Availability Zone deployment of the primary OB instance fails?

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

Answer: D

NEW QUESTION 2
- (Topic 1)
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: B

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

NEW QUESTION 3
- (Topic 1)
You are managing a legacy application Inside VPC with hard coded IP addresses in its configuration.
Which two mechanisms will allow the application to failover to new instances without the need for reconfiguration? Choose 2 answers

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

Answer: AD

NEW QUESTION 4
- (Topic 1)
You have set up Individual AWS accounts for each project. You have been asked to make sure your AWS Infrastructure costs do not exceed the budget set per project for each month.
Which of the following approaches can help ensure that you do not exceed the budget each month?

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

Answer: C

NEW QUESTION 5
- (Topic 1)
Your entire AWS infrastructure lives inside of one Amazon VPC You have an Infrastructure monitoring application running on an Amazon instance in Availability Zone (AZ) A of the region, and another application instance running in AZ B. The monitoring application needs to make use of ICMP ping to confirm network reachability of the instance hosting the application.
Can you configure the security groups for these instances to only allow the ICMP ping to pass from the monitoring instance to the application instance and nothing else'' If so how?

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

Answer: D

NEW QUESTION 6
- (Topic 1)
You have a web application leveraging an Elastic Load Balancer (ELB) In front of the web servers deployed using an Auto Scaling Group Your database is running

on Relational
Database Service (RDS) The application serves out technical articles and responses to them in general there are more views of an article than there are responses to the article. On occasion, an article on the site becomes extremely popular resulting in significant traffic Increases that causes the site to go down. What could you do to help alleviate the pressure on the infrastructure while maintaining availability during these events?
Choose 3 answers

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

Answer: ACE

NEW QUESTION 7
- (Topic 1)
You are tasked with the migration of a highly trafficked Node JS application to AWS In order to comply with organizational standards Chef recipes must be used to configure the application servers that host this application and to support application lifecycle events.
Which deployment option meets these requirements while minimizing administrative burden?

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

Answer: B

Explanation:
 Reference:
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deployment.html

NEW QUESTION 8
- (Topic 1)
You are running a database on an EC2 instance, with the data stored on Elastic Block Store (EBS) for persistence At times throughout the day, you are seeing large variance in the response times of the database queries Looking into the instance with the isolate command you see a lot of wait time on the disk volume that the database's data is stored on.
What two ways can you improve the performance of the database's storage while maintaining the current persistence of the data?
Choose 2 answers

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

Answer: AB

NEW QUESTION 9
- (Topic 1)
An application that you are managing has EC2 instances & Dynamo OB tables deployed to several AWS Regions In order to monitor the performance of the application globally, you would like to see two graphs 1) Avg CPU Utilization across all EC2 instances and 2) Number of Throttled Requests for all DynamoDB tables.
How can you accomplish this?

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

Answer: C

NEW QUESTION 10
- (Topic 1)
Which of the following requires a custom CloudWatch metric to monitor?

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

Answer: C

Explanation:
 Reference:
http://aws.amazon.com/cloudwatch/

NEW QUESTION 10

- (Topic 1)
The majority of your Infrastructure is on premises and you have a small footprint on AWS Your company has decided to roll out a new application that is heavily dependent on low latency connectivity to LOAP for authentication Your security policy requires minimal changes to the company's existing application user management processes.
What option would you implement to successfully launch this application1?

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

Answer: D

Explanation:
 Reference:
http://msdn.microsoft.com/en-us/library/azure/jj156090.aspx

NEW QUESTION 15
- (Topic 1)
Your company Is moving towards tracking web page users with a small tracking
Image loaded on each page Currently you are serving this image out of US-East, but are starting to get concerned about the time It takes to load the image for users on the west coast.
What are the two best ways to speed up serving this image?
Choose 2 answers

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

Answer: AD

NEW QUESTION 20
- (Topic 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 23
- (Topic 1)
You have identified network throughput as a bottleneck on your m1.small EC2 instance when uploading data Into Amazon S3 In the same region.
How do you remedy this situation?

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

Answer: B

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

NEW QUESTION 25
- (Topic 1)
You are using ElastiCache Memcached to store session state and cache database queries in your infrastructure. You notice in CloudWatch that Evictions and GetMisses are Doth very high.
What two actions could you take to rectify this?
Choose 2 answers

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

Answer: BD

NEW QUESTION 30
- (Topic 1)
An organization's security policy requires multiple copies of all critical data to be replicated across at least a primary and backup data center. The organization has decided to store some critical data on Amazon S3.
Which option should you implement to ensure this requirement is met?

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

Answer: D

NEW QUESTION 32
- (Topic 1)
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 37
- (Topic 2)
A user is accessing RDS from an application. The user has enabled the Multi AZ feature with the MS SQL RDS DB. During a planned outage how will AWS ensure that a switch from DB to a standby replica will not affect access to the application?

A. RDS will have an internal IP which will redirect all requests to the new DB
B. RDS uses DNS to switch over to stand by replica for seamless transition
C. The switch over changes Hardware so RDS does not need to worry about access
D. RDS will have both the DBs running independently and the user has to manually switch over

Answer: B

Explanation:

In the event of a planned or unplanned outage of a DB instance, Amazon RDS automatically switches to a standby replica in another Availability Zone if the user has enabled Multi AZ. The automatic failover mechanism simply changes the DNS record of the DB instance to point to the standby DB instance. As a result, the user will need to re-establish any existing connections to the DB instance. However, as the DNS is the same, the application can access DB seamlessly.

NEW QUESTION 42
- (Topic 2)
A user is planning to setup infrastructure on AWS for the Christmas sales. The user is planning to use Auto Scaling based on the schedule for proactive scaling. What advise would you give to the user?

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

Answer: C

Explanation:

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

NEW QUESTION 45
- (Topic 2)
A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created a public subnet CIDR (20.0.0.0/24. and VPN only subnets CIDR (20.0.1.0/24. along with the VPN gateway (vgw-12345. to connect to the user's data centre. Which of the below mentioned options is a valid entry for the main route table in this scenario?

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

Answer: D

Explanation:

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

route all internet traffic to the VPN gateway. Destination: 20.0.0.0/16 & Target: local (To allow local routing in VPC.

NEW QUESTION 47
- (Topic 2)
A user has setup an EBS backed instance and a CloudWatch alarm when the CPU utilization is more than 65%. The user has setup the alarm to watch it for 5 periods of 5 minutes each. The CPU utilization is 60% between 9 AM to 6 PM. The user has stopped the EC2 instance for 15 minutes between 11 AM to 11:15 AM. What will be the status of the alarm at 11:30 AM?

A. Alarm
B. OK
C. Insufficient Data
D. Error

Answer: B

Explanation:

Amazon CloudWatch alarm watches a single metric over a time period 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 state of the alarm will be OK for the whole day. When the user stops the instance for three periods the alarm may not receive the data

NEW QUESTION 52
- (Topic 2)
A user has created a VPC with CIDR 20.0.0.0/16. The user has created public and VPN only subnets along with hardware VPN access to connect to the user's datacenter. The user wants to make so that all traffic coming to the public subnet follows the organization's proxy policy. How can the user make this happen?

A. Setting up a NAT with the proxy protocol and configure that the public subnet receives traffic from NAT
B. Settin up a proxy policy in the internet gateway connected with the public subnet
C. It is not possible to setup the proxy policy for a public subnet
D. Setting the route table and security group of the public subnet which receives traffic from a virtual private gateway

Answer: D

Explanation:

The user can create subnets within a VPC. If the user wants to connect to VPC from his own data centre, he can setup public and VPN only subnets which uses hardware VPN access to connect with his data centre. When the user has configured this setup, 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. By default the internet traffic of the VPN subnet is routed to a virtual private gateway while the internet traffic of the public subnet is routed through the internet gateway. The user can set up the route and security group rules. These rules enable the traffic to come from the organization's network over the virtual private gateway to the public subnet to allow proxy settings on that public subnet.

NEW QUESTION 56
- (Topic 2)
An admin is planning to monitor the ELB. Which of the below mentioned services does not help the admin capture the monitoring information about the ELB activity?

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

Answer: B

Explanation:

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

NEW QUESTION 59
- (Topic 2)
An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

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

Answer: C

Explanation:

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

NEW QUESTION 60
- (Topic 2)
A user has created an ELB with three instances. How many security groups will ELB create by default?

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

Answer: C

Explanation:

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

NEW QUESTION 61
- (Topic 2)
A sys admin has created the below mentioned policy and applied to an S3 object named aws.jpg. The aws.jpg is inside a bucket named cloudacademy. What does this policy define?
"Statement": [{
"Sid": "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

Explanation:

A system admin can grant permission to the S3 objects or buckets to any user or make objects public using the bucket policy and user policy. Both use the JSON-based access policy language. Generally if the user is defining the ACL on the bucket, the objects in the bucket do not inherit it and vice a versa. The bucket policy can be defined at the bucket level which allows the objects as well as the bucket to be public with a single policy applied to that bucket. It cannot be applied at the object level.

NEW QUESTION 66
- (Topic 2)
A user has created a photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly.Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

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

Answer: D

Explanation:

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

NEW QUESTION 71
- (Topic 2)
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 73
- (Topic 2)
A user has configured Elastic Load Balancing by enabling a Secure Socket Layer (SSL. negotiation
configuration known as a Security Policy. Which of the below mentioned options is not part of this secure policy while negotiating the SSL connection between the user and the client?

A. SSL Protocols
B. Client Order Preference
C. SSL Ciphers
D. Server Order Preference

Answer: B

Explanation:

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

NEW QUESTION 78
- (Topic 2)
A user is checking the CloudWatch metrics from the AWS console. The user notices that the CloudWatch data is coming in UTC. The user wants to convert the data to a local time zone. How can the user perform this?

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

Answer: B

Explanation:

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

NEW QUESTION 83
- (Topic 2)
A user is planning to setup notifications on the RDS DB for a snapshot. Which of the below mentioned event categories is not supported by RDS for this snapshot source type?

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

Answer: A

Explanation:

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

NEW QUESTION 86
- (Topic 2)
A user has created a web application with Auto Scaling. The user is regularly monitoring the application and he observed that the traffic is highest on Thursday and Friday between 8 AM to 6 PM. What is the best solution to handle scaling in this case?

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

Answer: B

Explanation:

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

NEW QUESTION 90
- (Topic 2)
A root account owner has created an S3 bucket testmycloud. The account owner wants to allow everyone to upload the objects as well as enforce that the person who uploaded the object should manage the permission of those objects. Which is the easiest way to achieve this?

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

Answer: C

Explanation:

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

NEW QUESTION 93
- (Topic 2)
A user has configured ELB with two EBS backed EC2 instances. The user is trying to understand the DNS access and IP support for ELB. Which of the below mentioned statements may not help the user understand the IP mechanism supported by ELB?

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

Answer: D

Explanation:

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

NEW QUESTION 98
- (Topic 2)
An organization has added 3 of his AWS accounts to consolidated billing. One of the AWS accounts has
purchased a Reserved Instance (RI. of a small instance size in the US-East-1a zone. All other AWS accounts are running instances of a small size in the same zone. What will happen in this case for the RI pricing?

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

Answer: C

Explanation:

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

NEW QUESTION 99
- (Topic 2)
A sys admin is trying to understand the Auto Scaling activities. Which of the below mentioned processes is not performed by Auto Scaling?

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

Answer: A

Explanation:

There are two primary types of Auto Scaling processes: Launch and Terminate, which launch or terminat
instances, respectively. Some other actions performed by Auto Scaling are:
AddToLoadbalancer,
AlarmNotification, HealthCheck, AZRebalance, ReplaceUnHealthy, and ScheduledActions.

NEW QUESTION 101
- (Topic 2)
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 102
- (Topic 2)
A user has configured the Auto Scaling group with the minimum capacity as 3 and the maximum capacity as 5. When the user configures the AS group, how many instances will Auto Scaling launch?

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

Answer: C

NEW QUESTION 104
- (Topic 2)
A user has configured CloudWatch monitoring on an EBS backed EC2 instance. If the user has not attached any additional device, which of the below mentioned metrics will always show a 0 value?

A. DiskReadBytes
B. NetworkIn
C. NetworkOut
D. CPUUtilization

Answer: A

Explanation:

CloudWatch is used to monitor AWS as the well custom services. For EC2 when the user is monitoring the EC2 instances, it will capture the 7 Instance level and 3 system check parameters for the EC2 instance. Since this is an EBS backed instance, it will not have ephermal storage attached to it. Out of the 7 EC2 metrics, the 4 metrics DiskReadOps, DiskWriteOps, DiskReadBytes and DiskWriteBytes are disk related data and available only when there is ephermal storage attached to an instance. For an EBS backed instance without any additional device, this data will be 0.

NEW QUESTION 106
- (Topic 2)
A user is trying to aggregate all the CloudWatch metric data of the last 1 week. Which of the below mentioned statistics is not available for the user as a part of data aggregation?

A. Aggregate
B. Sum
C. Sample data
D. Average

Answer: A

Explanation:

Amazon CloudWatch is basically a metrics repository. Either the user can send the custom data or an AWS product can put metrics into the repository, and the user can retrieve the statistics based on those metrics. The statistics are metric data aggregations over specified periods of time. Aggregations are made using the namespace, metric name, dimensions, and the data point unit of measure, within the time period that is specified by the user. CloudWatch supports Sum, Min, Max, Sample Data and Average statistics aggregation.

NEW QUESTION 107
- (Topic 2)
A user has configured an Auto Scaling group with ELB. The user has enabled detailed CloudWatch monitoring on Elastic Load balancing. Which of the below mentioned statements will help the user understand this functionality better?

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

Answer: A

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. Elastic Load Balancing includes 10 metrics and 2 dimensions, and sends data to CloudWatch every minute. This does not cost extra.

NEW QUESTION 110
- (Topic 2)
An organization has created 50 IAM users. The organization wants that each user can change their password but cannot change their access keys. How can the organization achieve this?

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

Answer: D

Explanation:

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

NEW QUESTION 113
- (Topic 2)
A user is trying to delete an Auto Scaling group from CLI. Which of the below mentioned steps are to be performed by the user?

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

Answer: C

Explanation:

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

NEW QUESTION 117
- (Topic 2)
An organization is using cost allocation tags to find the cost distribution of different departments and projects. One of the instances has two separate tags with the key/ value as "InstanceName/HR", "CostCenter/HR". What will AWS do in this case?

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

Answer: D

Explanation:

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

NEW QUESTION 119
- (Topic 2)
A sys admin is trying to understand EBS snapshots. Which of the below mentioned statements will not be
useful to the admin to understand the concepts about a snapshot?

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

Answer: A

Explanation:

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

NEW QUESTION 121
- (Topic 3)
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 126
- (Topic 3)
A user has deployed an application on an EBS backed EC2 instance. For a better performance of application, it requires dedicated EC2 to EBS traffic. How can the user achieve this?

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

Answer: D

Explanation:

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

NEW QUESTION 129
- (Topic 3)
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 week

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 131
- (Topic 3)
A sys admin has enabled logging on ELB. Which of the below mentioned fields will not be a part of the log file name?

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

Answer: B

Explanation:

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

NEW QUESTION 136
- (Topic 3)
A user has created a VPC with CIDR 20.0.0.0/24. The user has used all the IPs of CIDR and wants to increase the size of the VPC. The user has two subnets: public (20.0.0.0/28. and private (20.0.1.0/28.. How can the user change the size of the VPC?

A. The user can delete all the instances of the subne
B. Change the size of the subnets to 20.0.0.0/32 and 20.0.1.0/32, 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 138
- (Topic 3)
A user has launched 5 instances in EC2-CLASSIC and attached 5 elastic IPs to the five different instances in the US East region. The user is creating a VPC in the same region. The user wants to assign an elastic IP to the VPC instance. How can the user achieve this?

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

Answer: D

Explanation:
 Section: (none)
A Virtual Private Cloud (VPC. is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. A user can have 5 IP addresses per region with EC2 Classic. The user can have 5 separate IPs with VPC in the same region as it has a separate limit than EC2 Classic.

NEW QUESTION 139
- (Topic 3)
A user has launched two EBS backed EC2 instances in the US-East-1a region. The user wants to change the zone of one of the instances. How can the user change it?

A. The zone can only be modified using the AWS CLI
B. It is not possible to change the zone of an instance after it is launched
C. Stop one of the instances and change the availability zone
D. From the AWS EC2 console, select the Actions - > Change zones and specify the new zone

Answer: B

Explanation:

With AWS EC2, when a user is launching an instance he can select the availability zone (AZ. at the time of launch. If the zone is not selected, AWS selects it on behalf of the user. Once the instance is launched, the user cannot change the zone of that instance unless he creates an AMI of that instance and launches a new instance from it.

NEW QUESTION 141
- (Topic 3)
A user is sending the data to CloudWatch using the CloudWatch API. The user is sending data 90 minutes in the future. What will CloudWatch do in this case?

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

Answer: A

Explanation:

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

NEW QUESTION 145
- (Topic 3)
A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. Which of the
below mentioned statements is true with respect to this scenario?

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

Answer: D

Explanation:

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

connect to the instance from the internet he should create an internet gateway and assign an elastic IP to instance.

NEW QUESTION 147
- (Topic 3)
An organization has configured two single availability zones. The Auto Scaling groups are configured in
separate zones. The user wants to merge the groups such that one group spans across multiple zones. How can the user configure this?

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

Answer: B

Explanation:

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

NEW QUESTION 151
- (Topic 3)
A user has setup a CloudWatch alarm on the EC2 instance for CPU utilization. The user has setup to receive a notification on email when the CPU utilization is higher than 60%. The user is running a virus scan on the same instance at a particular time. The user wants to avoid receiving an email at this time. What should the user do?

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

Answer: B

Explanation:

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

NEW QUESTION 155
- (Topic 3)
In AWS, which security aspects are the customer's responsibility? Choose 4 answers

A. Controlling physical access to compute resources
B. Patch management on the EC2 instance s operating system
C. Encryption of EBS (Elastic Block Storage) volumes
D. Life-cycle management of IAM credentials
E. Decommissioning storage devices
F. Security Group and ACL (Access Control List) settings

Answer: BCEF

NEW QUESTION 158
- (Topic 3)
A user is using the AWS SQS to decouple the services. Which of the below mentioned operations is not
supported by SQS?

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

Answer: D

Explanation:

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

NEW QUESTION 159
- (Topic 3)
A user is having data generated randomly based on a certain event. The user wants to upload that data to CloudWatch. It may happen that event may not have data generated for some period due to andomness. Which of the below mentioned options is a recommended option for this case?

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

Answer: C

Explanation:

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

NEW QUESTION 162
- (Topic 3)
Which services allow the customer to retain run administrative privileges or the underlying EC2 instances? Choose 2 answers

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

Answer: AB

NEW QUESTION 164
- (Topic 3)
A root account owner has given full access of his S3 bucket to one of the IAM users using the bucket ACL. When the IAM user logs in to the S3 console, which actions can he perform?

A. He can just view the content of the bucket
B. He can do all the operations on the bucket
C. It is not possible to give access to an IAM user using ACL
D. The IAM user can perform all operations on the bucket using only API/SDK

Answer: C

Explanation:

Each AWS S3 bucket and object has an ACL (Access Control List. associated with it. An ACL is a list of grants identifying the grantee and the permission granted. The user can use ACLs to grant basic read/write permissions to other AWS accounts. ACLs use an Amazon S3–specific XML schema. The user cannot grant permissions to other users (IAM users. in his account.

NEW QUESTION 168
- (Topic 3)
A user has configured Auto Scaling with 3 instances. The user had created a new AMI after updating one of the instances. If the user wants to terminate two specific instances to ensure that Auto Scaling launches an instances with the new launch configuration, which command should he run?

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

Answer: D

Explanation:

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

NEW QUESTION 172
- (Topic 3)
George has shared an EC2 AMI created in the US East region from his AWS account with Stefano. George copies the same AMI to the US West region. Can Stefano access the copied AMI of George's account from the US West region?

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

Answer: A

Explanation:

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

NEW QUESTION 173
- (Topic 3)
An organization (Account ID 123412341234. has attached the below mentioned IAM policy to a user. What does this policy statement entitle the user to perform?
{
"Version": "2012-10-17",
"Statement": [{
"Sid": "AllowUsersAllActionsForCredentials",
"Effect": "Allow",
"Action": [
"iam:*LoginProfile",
"iam:*AccessKey*",
"iam:*SigningCertificate*"
],
"Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
}]
}

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

Answer: D

Explanation:

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

NEW QUESTION 176
- (Topic 3)
An organization has applied the below mentioned policy on an IAM group which has selected the IAM users. What entitlements do the IAM users avail with this policy?
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "*",
"Resource": "*"
}
]
}

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

Answer: C

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The IAM group allows the organization to specify permissions for a collection of users. With the below mentioned policy, it will allow the group full access (Admin. to all AWS services.
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "*",
"Resource": "*"

```
}
]
}
```

NEW QUESTION 179
- (Topic 3)
The compliance department within your multi-national organization requires that all data for your customers that reside in the European Union (EU) must not leave the EU and also
data for customers that reside in the US must not leave the US without explicit authorization.
What must you do to comply with this requirement for a web based profile management application running on EC2?

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

Answer: C

NEW QUESTION 183
- (Topic 3)
A user is planning to use AWS services for his web application. If the user is trying to set up his own billing management system for AWS, how can he configure it?

A. Set up programmatic billing acces
B. Download and parse the bill as per the requirement
C. It is not possible for the user to create his own billing management service with AWS
D. Enable the AWS CloudWatch alarm which will provide APIs to download the alarm data
E. Use AWS billing APIs to download the usage report of each service from the AWS billing console

Answer: A

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. AWS will upload the bill to the bucket every few hours and the user can download the bill CSV from the bucket, parse itand create a billing system as per the requirement.

NEW QUESTION 187
- (Topic 3)
A user has created an EBS volume of 10 GB and attached it to a running instance. The user is trying to access EBS for first time. Which of the below mentioned options is the correct statement with respect to a first time EBS access?

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

Answer: B

Explanation:

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

NEW QUESTION 189
- (Topic 3)
An organization has configured Auto Scaling for hosting their application. The system admin wants to
understand the Auto Scaling health check process. If the instance is unhealthy, Auto Scaling launches an
instance and terminates the unhealthy instance. What is the order execution?

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

Answer: D

Explanation:

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

NEW QUESTION 194
- (Topic 3)
A user has configured Auto Scaling with the minimum capacity as 2 and the desired capacity as 2. The user is trying to terminate one of the existing instance with the command:
as-terminate-instance-in-auto-scaling-group<Instance ID> --decrement-desired-capacity
What will Auto Scaling do in this scenario?

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

Answer: D

Explanation:

The Auto Scaling command as-terminate-instance-in-auto-scaling-group <Instance ID> will terminate the specific instance ID. The user is required to specify the parameter 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 199
- (Topic 3)
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 204
- (Topic 3)
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 207
- (Topic 3)
An application you maintain consists of multiple EC2 instances in a default tenancy VPC. This application has undergone an internal audit and has been determined to require dedicated hardware for one instance. Your compliance team has given you a week to move this instance to single-tenant hardware.
Which process will have minimal impact on your application while complying with this requirement?

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

Answer: A

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

NEW QUESTION 210
- (Topic 3)
A user has created a VPC with the public subnet. The user has created a security group for that VPC. Which of the below mentioned statements is true when a security group is created?

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

Answer: C

Explanation:

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

NEW QUESTION 211
- (Topic 3)
A user has enabled versioning on an S3 bucket. The user is using server side encryption for data at rest. If the user is supplying his own keys for encryption (SSE-C., what is recommended to the user for the purpose of security?

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

Answer: D

Explanation:

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

NEW QUESTION 215
- (Topic 3)
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 220
- (Topic 3)
A user has a weighing plant. The user measures the weight of some goods every 5 minutes and sends data to AWS CloudWatch for monitoring and tracking. Which of the below mentioned parameters is mandatory for the user to include in the request list?

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

Answer: B

Explanation:

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

NEW QUESTION 225
- (Topic 3)
A user has setup a custom application which generates a number in decimals. The user wants to track that number and setup the alarm whenever the number is above a certain limit. The application is sending the data to CloudWatch at regular intervals for this purpose. Which of the below mentioned statements is not true with respect to the above scenario?

A. The user can get the aggregate data of the numbers generated over a minute and send it to CloudWatch
B. The user has to supply the timezone with each data point
C. CloudWatch will not truncate the number until it has an exponent larger than 126 (i.

D. (1 x 10^126.
E. The user can create a file in the JSON format with the metric name and value and supply it to CloudWatch

Answer: B

NEW QUESTION 229
- (Topic 3)
Amazon EBS snapshots have which of the following two characteristics? (Choose 2.) Choose 2 answers

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

Answer: AD

NEW QUESTION 230
- (Topic 3)
A user has scheduled the maintenance window of an RDS DB on Monday at 3 AM. Which of the below mentioned events may force to take the DB instance offline during the maintenance window?

A. Enabling Read Replica
B. Making the DB Multi AZ
C. DB password change
D. Security patching

Answer: D

Explanation:

Amazon RDS performs maintenance on the DB instance during a user-definable maintenance window. The system may be offline or experience lower performance during that window. The only maintenance events that may require RDS to make the DB instance offline are: Scaling compute operations Software patching. Required software patching is automatically scheduled only for patches that are security and durability related. Such patching occurs infrequently (typically once every few months. and seldom requires more than a fraction of the maintenance window.

NEW QUESTION 232
- (Topic 3)
A user has launched an EC2 instance from an instance store backed AMI. The infrastructure team wants to create an AMI from the running instance. Which of the below mentioned credentials is not required while creating the AMI?

A. AWS account ID
B. X.509 certificate and private key
C. AWS login ID to login to the console
D. Access key and secret access key

Answer: C

Explanation:

When the user has launched an EC2 instance from an instance store backed AMI and the admin team wants to create an AMI from it, the user needs to setup the AWS AMI or the API tools first. Once the tool is setup the user will need the following credentials:
AWS account ID;
AWS access and secret access key;
X.509 certificate with private key.

NEW QUESTION 236
- (Topic 3)
A user had aggregated the CloudWatch metric data on the AMI ID. The user observed some abnormal
behaviour of the CPU utilization metric while viewing the last 2 weeks of data. The user wants to share that data with his manager. How can the user achieve this easily with the AWS console?

A. The user can use the copy URL functionality of CloudWatch to share the exact details
B. The user can use the export data option from the CloudWatch console to export the current data point
C. The user has to find the period and data and provide all the aggregation information to the manager
D. The user can use the CloudWatch data copy functionality to copy the current data points

Answer: A

Explanation:

Amazon CloudWatch provides the functionality to graph the metric data generated either by the AWS services or the custom metric to make it easier for the user to analyse. The console provides the option to save the URL or bookmark it so that it can be used in the future by typing the same URL. The Copy URL functionality is available under the console when the user selects any metric to view.

NEW QUESTION 238
- (Topic 3)
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 Amazon AWS-SysOps : Practice Test
to launch, Cloudformation will rollback all the changes and terminate or delete all the created services.

NEW QUESTION 240
- (Topic 3)
A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console?

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

Answer: D

Explanation:

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

NEW QUESTION 243
- (Topic 3)
A user is trying to send custom metrics to CloudWatch using the PutMetricData APIs. Which of the below mentioned points should the user needs to take care while sending the data to CloudWatch?

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

Answer: A

Explanation:

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

NEW QUESTION 245
- (Topic 3)
A user runs the command "dd if=/dev/xvdf of=/dev/null bs=1M" on an EBS volume created from a snapshot and attached to a Linux instance. Which of the below mentioned activities is the user performing with the step given above?

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

Answer: A

Explanation:

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

NEW QUESTION 250
- (Topic 3)
A sys admin is trying to understand the sticky session algorithm. Please select the correct sequence of steps, both when the cookie is present and when it is not, to help the admin understand the implementation of the sticky session:
ELB inserts the cookie in the response ELB chooses the instance based on the load balancing algorithm Check the cookie in the service request The cookie is

found in the request The cookie is not found in the request

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

Answer: C

Explanation:

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

NEW QUESTION 252
- (Topic 3)
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 257
- (Topic 3)
A user has configured ELB with SSL using a security policy for secure negotiation between the client and load balancer. The ELB security policy supports various ciphers. Which of the below mentioned options helps identify the matching cipher at the client side to the ELB cipher list when client is requesting ELB DNS over SSL?

A. Cipher Protocol
B. Client Configuration Preference
C. Server Order Preference
D. Load Balancer Preference

Answer: C

Explanation:

Elastic Load Balancing uses a Secure Socket Layer (SSL. negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. When client is requesting ELB DNS over SSL and if the load balancer is configured to support the Server Order Preference, then the load balancer gets to select the first cipher in its list that matches any one of the ciphers in the client's list. Server Order Preference ensures that the load balancer determines which cipher is used for the SSL connection.

NEW QUESTION 260
- (Topic 3)
A user has configured an ELB to distribute the traffic among multiple instances. The user instances are facing some issues due to the back-end servers. Which of the below mentioned CloudWatch metrics helps the user understand the issue with the instances?

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

Answer: D

Explanation:

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

NEW QUESTION 262
- (Topic 3)

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 the internet?

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 265
- (Topic 3)
An AWS account wants to be part of the consolidated billing of his organization's payee account. How can the owner of that account achieve this?

A. The payee account has to request AWS support to link the other accounts with his account
B. The owner of the linked account should add the payee account to his master account list from the billing console
C. The payee account will send a request to the linked account to be a part of consolidated billing
D. The owner of the linked account requests the payee account to add his account to consolidated billing

Answer: C

Explanation:

AWS consolidated billing enables the organization to consolidate payments for multiple Amazon Web Services (AWS. accounts within a single organization by making a single paying account. To add a particular account (linked. to the master (payee. account, the payee account has to request the linked account to join consolidated billing. Once the linked account accepts the request henceforth all charges incurred by the linked account will be paid by the payee account.

NEW QUESTION 268
- (Topic 3)
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 metadat
B. Query the appropriate Amazon CloudWatch metri
C. Query the local instance userdat
D. Use ipconfig or ifconfig comman

Answer: B

NEW QUESTION 272
- (Topic 3)
A user has launched multiple EC2 instances for the purpose of development and testing in the same region. The user wants to find the separate cost for the production and development instances. How can the user find the cost distribution?

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

Answer: D

Explanation:

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

NEW QUESTION 273
- (Topic 3)
A user has configured ELB with SSL using a security policy for secure negotiation between the client and load balancer. Which of the below mentioned SSL protocols is not supported by the security policy?

A. TLS 1.3
B. TLS 1.2
C. SSL 2.0
D. SSL 3.0

Answer: A

Explanation:

Elastic Load Balancing uses a Secure Socket Layer (SSL. negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. Elastic Load Balancing supports the following versions of the SSL protocol: TLS 1.2 TLS 1.1 TLS 1.0 SSL 3.0

SSL 2.0

NEW QUESTION 276
- (Topic 3)
A user has enabled session stickiness with ELB. The user does not want ELB to manage the cookie; instead he wants the application to manage the cookie. What will happen when the server instance, which is bound to a cookie, crashes?

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

Answer: B

Explanation:

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

NEW QUESTION 280
- (Topic 3)
A user has launched an EC2 instance. The instance got terminated as soon as it was launched. Which of the below mentioned options is not a possible reason for this?

A. The user account has reached the maximum EC2 instance limit
B. The snapshot is corrupt
C. The AMI is missin
D. It is the required part
E. The user account has reached the maximum volume limit

Answer: A

Explanation:

When the user account has reached the maximum number of EC2 instances, it will not be allowed to launch an instance. AWS will throw an 'InstanceLimitExceeded' error. For all other reasons, such as "AMI is missing part", "Corrupt Snapshot" or "Volume limit has reached" it will launch an EC2 instance and then terminate it.

NEW QUESTION 282
- (Topic 3)
A user is configuring the Multi AZ feature of an RDS DB. The user came to know that this RDS DB does not use the AWS technology, but uses server mirroring to achieve HA. Which DB is the user using right now?

A. My SQL
B. Oracle
C. MS SQL
D. PostgreSQL

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. Multi AZ deployments for Oracle, PostgreSQL, and MySQL DB instances use Amazon technology, while SQL Server (MS SQL. DB instances use SQL Server Mirroring.

NEW QUESTION 286
- (Topic 3)
A user has created an application which will be hosted on EC2. The application makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK to connect with from the EC2 instance. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

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

Answer: A

Explanation:

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

NEW QUESTION 290
- (Topic 3)
A user is using the AWS EC2. The user wants to make so that when there is an issue in the EC2 server, such as instance status failed, it should start a new instance in the user's private cloud. Which AWS service helps to achieve this automation?

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

Answer: D

Explanation:

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

NEW QUESTION 293
- (Topic 3)
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 295
- (Topic 3)
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 296
- (Topic 3)
A sys admin is using server side encryption with AWS S3. Which of the below mentioned statements helps the user understand the S3 encryption functionality?

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

Answer: A

Explanation:

AWS S3 supports client side or server side encryption to encrypt all data at rest. The server side encryption can either have the S3 supplied AES-256 encryption key or the user can send the key along with each API call to supply his own encryption key. The encryption with the user supplied key (SSE-C. does not work with the AWS console. The S3 does not store the keys and the user has to send a key with each request. The SSE-C works when the user has enabled versioning.

NEW QUESTION 297
- (Topic 3)
An AWS root account owner is trying to create a policy to access RDS. Which of the below mentioned
statements is true with respect to the above information?

A. Create a policy which allows the users to access RDS and apply it to the RDS instances

B. The user cannot access the RDS database if he is not assigned the correct IAM policy
C. The root account owner should create a policy for the IAM user and give him access to the RDS services
D. The policy should be created for the user and provide access for RDS

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 account owner wants to create a policy for RDS, the owner has to create an IAM user and define the policy which entitles the IAM user with various RDS services such as Launch Instance, Manage security group, Manage parameter group etc.

NEW QUESTION 299

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

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

Answer: B

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

NEW QUESTION 302

A user is trying to setup a scheduled scaling activity using Auto Scaling. The user wants to setup the recurring schedule. Which of the below mentioned parameters is not required in this case?

A. Maximum size
B. Auto Scaling group name
C. End time
D. Recurrence value

Answer: A

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. If the user is setting a recurring event, it is required that the user specifies the Recurrence value (in a cron format., end time (not compulsory but recurrence will stop after this. and the Auto Scaling group for which the scaling activity is to be scheduled.

NEW QUESTION 306
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-SysOps Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-SysOps Product From:

https://www.2passeasy.com/dumps/AWS-SysOps/

Money Back Guarantee

AWS-SysOps Practice Exam Features:

* AWS-SysOps Questions and Answers Updated Frequently

* AWS-SysOps Practice Questions Verified by Expert Senior Certified Staff

* AWS-SysOps Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-SysOps Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year