Amazon

Exam Questions AWS-SysOps

Amazon AWS Certified SysOps Administrator - Associate

NEW QUESTION 1
- (Topic 1)
Your team Is excited about the use of AWS because now they have access to programmable Infrastructure" You have been asked to manage your AWS infrastructure In a manner similar to the way you might manage application code You want to be able to deploy exact copies of different versions of your infrastructure, stage changes into different environments, revert back to previous versions, and identify what versions are running at any particular time (development test QA. production).
Which approach addresses this requirement?

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

Answer: B

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

NEW QUESTION 2
- (Topic 1)
A media company produces new video files on-premises every day with a total size of around 100GBS after compression All files have a size of 1 -2 GB and need to be uploaded to Amazon S3 every night in a fixed time window between 3am and 5am Current upload takes almost 3 hours, although less than half of the available bandwidth is used.
What step(s) would ensure that the file uploads are able to complete in the allotted time window?

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

Answer: B

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

NEW QUESTION 3
- (Topic 1)
You are tasked with setting up a cluster of EC2 Instances for a NoSQL database. The database requires random read IO disk performance up to a 100,000 IOPS at 4KB block side per node.
Which of the following EC2 instances will perform the best for this workload?

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

Answer: C

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

NEW QUESTION 4
- (Topic 1)
Which of the following statements about this S3 bucket policy is true?

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

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

Answer: B

NEW QUESTION 5
- (Topic 1)
A customer has a web application that uses cookie Based sessions to track logged in users It Is deployed on AWS using ELB and Auto Scaling The customer observes that when load increases. Auto Scaling launches new Instances but the load on the easting Instances does not decrease, causing all existing users to have a sluggish experience.
Which two answer choices independently describe a behavior that could be the cause of the sluggish user experience? Choose 2 answers

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

Answer: BD

NEW QUESTION 6
- (Topic 1)
When preparing for a compliance assessment of your system built inside of AWS. what are three best-practices for you to prepare for an audit?
Choose 3 answers

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

Answer: ABD

NEW QUESTION 7
- (Topic 1)
You need to design a VPC for a web-application consisting of an Elastic Load Balancer (ELB). a fleet of web/application servers, and an RDS database The entire Infrastructure must be distributed over 2 availability zones.
Which VPC configuration works while assuring the database is not available from the Internet?

A. One public subnet for ELB one public subnet for the web-servers, and one private subnet for the database
B. One public subnet for ELB two private subnets for the web-servers, two private subnets for RDS
C. Two public subnets for ELB two private subnets for the web-servers and two private subnets for RDS
D. Two public subnets for ELB two public subnets for the web-servers, and two public subnets for RDS

Answer: A

NEW QUESTION 8

- (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 9
- (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 10
- (Topic 1)
You run a web application where web servers on EC2 Instances are In an Auto Scaling group Monitoring over the last 6 months shows that 6 web servers are necessary to handle the minimum load During the day up to 12 servers are needed Five to six days per year, the number of web servers required might go up to 15.
What would you recommend to minimize costs while being able to provide hill availability?

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

Answer: B

NEW QUESTION 10
- (Topic 1)
You are currently hosting multiple applications in a VPC and have logged numerous port scans coming in from a specific IP address block. Your security team has requested that all access from the offending IP address block be denied for the next 24 hours.
Which of the following is the best method to quickly and temporarily deny access from the specified IP address block?

A. Create an AD policy to modify Windows Firewall settings on all hosts in the VPC to deny access from the IP address block
B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP address block
C. Add a rule to all of the VPC 5 Security Groups to deny access from the IP address block
D. Modify the Windows Firewall settings on all Amazon Machine Images (AMIs) that your organization uses in that VPC to deny access from the IP address block

Answer: B

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

NEW QUESTION 13
- (Topic 1)
Which of the following are characteristics of Amazon VPC subnets?
Choose 2 answers

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

Answer: CE

NEW QUESTION 17
- (Topic 1)
You are creating an Auto Scaling group whose Instances need to insert a custom metric into CloudWatch.
Which method would be the best way to authenticate your CloudWatch PUT request?

A. Create an IAM role with the Put MetricData permission and modify the Auto Scaling launch configuration to launch instances in that role
B. Create an IAM user with the PutMetricData permission and modify the Auto Scaling launch configuration to inject the userscredentials into the instance User Data

C. Modify the appropriate Cloud Watch metric policies to allow the Put MetricData permission to instances from the Auto Scaling group
D. Create an IAM user with the PutMetricData permission and put the credentials in a private repository and have applications on the server pull the credentials as needed

Answer: A

NEW QUESTION 20
- (Topic 1)
When an EC2 instance that is backed by an S3-based AMI Is terminated, what happens to the data on me root volume?

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

Answer: C

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

NEW QUESTION 22
- (Topic 1)
You have an Auto Scaling group associated with an Elastic Load Balancer (ELB). You have noticed that instances launched via the Auto Scaling group are being marked unhealthy due to an ELB health check, but these unhealthy instances are not being terminated
What do you need to do to ensure trial instances marked unhealthy by the ELB will be terminated and replaced?

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

Answer: B

Explanation:
 Reference:
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-add-elb-healthcheck.html
Add an Elastic Load Balancing Health Check to your Auto Scaling Group
By default, an Auto Scaling group periodically reviews the results of EC2 instance status to determine the health state of each instance. However, if you have associated your Auto Scaling group with an Elastic Load Balancing load balancer, you can choose to use the Elastic Load Balancing health check. In this case, Auto Scaling determines the health status of your instances by checking the results of both the EC2 instance status check and the Elastic Load Balancing instance health check.
For information about EC2 instance status checks, see.Monitor Instances With Status Checks.in the.Amazon EC2 User Guide for Linux Instances. For information about Elastic Load Balancing health checks, see.Health Check.in the.Elastic Load Balancing Developer Guide.
This topic shows you how to add an Elastic Load Balancing health check to your Auto Scaling group, assuming that you have created a load balancer and have registered the load balancer with your Auto Scaling group. If you have not registered the load balancer with your Auto Scaling group, see.Set Up a Scaled and Load-Balanced Application.
Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action.DescribeInstanceStatus.return any state other than.running, the system status shows.impaired, or the calls to Elastic Load Balancing action.DescribeInstanceHealth.returns.OutOfService.in the instance state field.
If there are multiple load balancers associated with your Auto Scaling group, Auto Scaling checks the health state of your EC2 instances by making health check calls to each load balancer. For each call, if the Elastic Load Balancing action returns any state other than.InService, the instance is marked as unhealthy. After Auto Scaling marks an instance as unhealthy, it remains in that state, even if subsequent calls from other load balancers return an.InService.state for the same instance.

NEW QUESTION 27
- (Topic 1)
You receive a frantic call from a new DBA who accidentally dropped a table containing all your customers.
Which Amazon RDS feature will allow you to reliably restore your database to within 5 minutes of when the mistake was made?

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

Answer: D

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

NEW QUESTION 28
- (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 33
- (Topic 1)
You have been asked to automate many routine systems administrator backup and recovery activities. Your current plan is to leverage AWS-managed solutions as much as possible and automate the rest with the AWS CLI and scripts.
Which task would be best accomplished with a script?

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

Answer: A

NEW QUESTION 37
- (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 39
- (Topic 1)
You have a Linux EC2 web server instance running inside a VPC The instance is In a public subnet and has an EIP associated with it so you can connect to It over the Internet via HTTP or SSH The instance was also fully accessible when you last logged in via SSH. and was also serving web requests on port 80.
Now you are not able to SSH into the host nor does it respond to web requests on port 80 that were working fine last time you checked You have double-checked that all networking configuration parameters (security groups route tables. IGW'EIP. NACLs etc) are properly configured {and you haven't made any changes to those anyway since you were last able to reach the Instance). You look at the EC2 console and notice that system status check shows "impaired."
Which should be your next step in troubleshooting and attempting to get the instance back to a healthy state so that you can log in again?

A. Stop and start the instance so that it will be able to be redeployed on a healthy host system that most likely will fix the "impaired" system status
B. Reboot your instance so that the operating system will have a chance to boot in a clean healthy state that most likely will fix the 'impaired" system status
C. Add another dynamic private IP address to me instance and try to connect via mat new path, since the networking stack of the OS may be locked up causing the "impaired" system statu
D. Add another Elastic Network Interface to the instance and try to connect via that new path since the networking stack of the OS may be locked up causing the "impaired" system status
E. un-map and then re-map the EIP to the instance, since the IGWVNAT gateway may not be working properly, causing the "impaired" system status

Answer: A

NEW QUESTION 40
- (Topic 2)
A user is trying to understand the ACL and policy for an S3 bucket. Which of the below mentioned policy
permissions is equivalent to the WRITE ACL on a bucket?

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

Answer: D

Explanation:

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

NEW QUESTION 45
- (Topic 2)
An application is generating a log file every 5 minutes. The log file is not critical but may be required only for verification in case of some major issue. The file should be accessible over the internet whenever required. Which of the below mentioned options is a best possible storage solution for it?

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

Answer: D

Explanation:

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

NEW QUESTION 49
- (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 54
- (Topic 2)
A system admin is managing buckets, objects and folders with AWS S3. Which of the below mentioned statements is true and should be taken in consideration by the sysadmin?

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

Answer: A

Explanation:

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

NEW QUESTION 57
- (Topic 2)
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 61
- (Topic 2)
A user has created a photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly.Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

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

Answer: D

Explanation:

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

NEW QUESTION 65
- (Topic 2)
A user is planning to evaluate AWS for their internal use. The user does not want to incur any charge on his account during the evaluation. Which of the below mentioned AWS services would incur a charge if used?

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

Answer: D

Explanation:

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

NEW QUESTION 66
- (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 69
- (Topic 2)
A user has launched 10 instances from the same AMI ID using Auto Scaling. The user is trying to see the
average CPU utilization across all instances of the last 2 weeks under the CloudWatch console. How can the user achieve this?

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

Answer: B

Explanation:

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

NEW QUESTION 73
- (Topic 2)
A user has created a queue named "myqueue" in US-East region with AWS SQS. The user's AWS account ID is 123456789012. If the user wants to perform some action on this queue, which of the below Queue URL should he use?

A. http://sqs.us-east-1.amazonaws.com/123456789012/myqueue
B. http://sqs.amazonaws.com/123456789012/myqueue
C. http://sq
D. 123456789012.us-east-1.amazonaws.com/myqueue
E. http:// 123456789012.sq
F. us-east-1.amazonaws.com/myqueue

Answer: A

Explanation:

When creating a new queue in SQS, the user must provide a queue name that is unique within the scope of all queues of user's account. If the user creates queues using both the latest WSDL and a previous version, he will have a single namespace for all his queues. Amazon SQS assigns each queue created by user an identifier called a queue URL, which includes the queue name and other components that Amazon SQS determines. Whenever the user wants to perform an

action on a queue, he must provide its queue URL. The queue URL for the account id 123456789012 & queue name "myqueue" in US-East-1 region will be http://sqs.us-east-1.amazonaws.com/123456789012/myqueue.

NEW QUESTION 78
- (Topic 2)
A user has launched an ELB which has 5 instances registered with it. The user deletes the
ELB by mistake. What will happen to the instances?

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

Answer: D

Explanation:

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

NEW QUESTION 82
- (Topic 2)
An organization has setup consolidated billing with 3 different AWS accounts. Which of the below mentioned advantages will organization receive in terms of the AWS pricing?

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

Answer: B

Explanation:

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

NEW QUESTION 84
- (Topic 2)
A user has setup connection draining with ELB to allow in-flight requests to continue while the instance is being deregistered through Auto Scaling. If the user has not specified the draining time, how long will ELB allow inflight requests traffic to continue?

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

Answer: C

Explanation:

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

NEW QUESTION 85
- (Topic 2)
A customer is using AWS for Dev and Test. The customer wants to setup the Dev environment with
Cloudformation. Which of the below mentioned steps are not required while using Cloudformation?

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

Answer: B

Explanation:

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

NEW QUESTION 88
- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 89
- (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 93
- (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 97
- (Topic 2)
A user has enabled detailed CloudWatch metric monitoring on an Auto Scaling group. Which of the below mentioned metrics will help the user identify the total number of instances in an Auto Scaling group cluding pending, terminating and running instances?

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

Answer: A

Explanation:

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

NEW QUESTION 99
- (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 104
- (Topic 2)
A user has launched an EBS backed instance. The user started the instance at 9 AM in the morning. Between 9 AM to 10 AM, the user is testing some script. Thus, he stopped the instance twice and restarted it. In the same hour the user rebooted the instance once. For how many instance hours will AWS charge the user?

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

Answer: A

Explanation:

A user can stop/start or reboot an EC2 instance using the AWS console, the Amazon EC2 CLI or the Amazon EC2 API. Rebooting an instance is equivalent to rebooting an operating system. When the instance is rebooted AWS will not charge the user for the extra hours. In case the user stops the instance, AWS does not charge the running cost but charges only the EBS storage cost. If the user starts and stops the instance multiple times in a single hour, AWS will charge the user for every start and stop. In this case, since the instance was rebooted twice, it will cost the user for 3 instance hours.

NEW QUESTION 109
- (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 110
- (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 115
- (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 116
- (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 121
- (Topic 2)
You are building an online store on AWS that uses SQS to process your customer orders. Your backend system needs those messages in the same sequence the customer orders have been put in. How can you achieve that?

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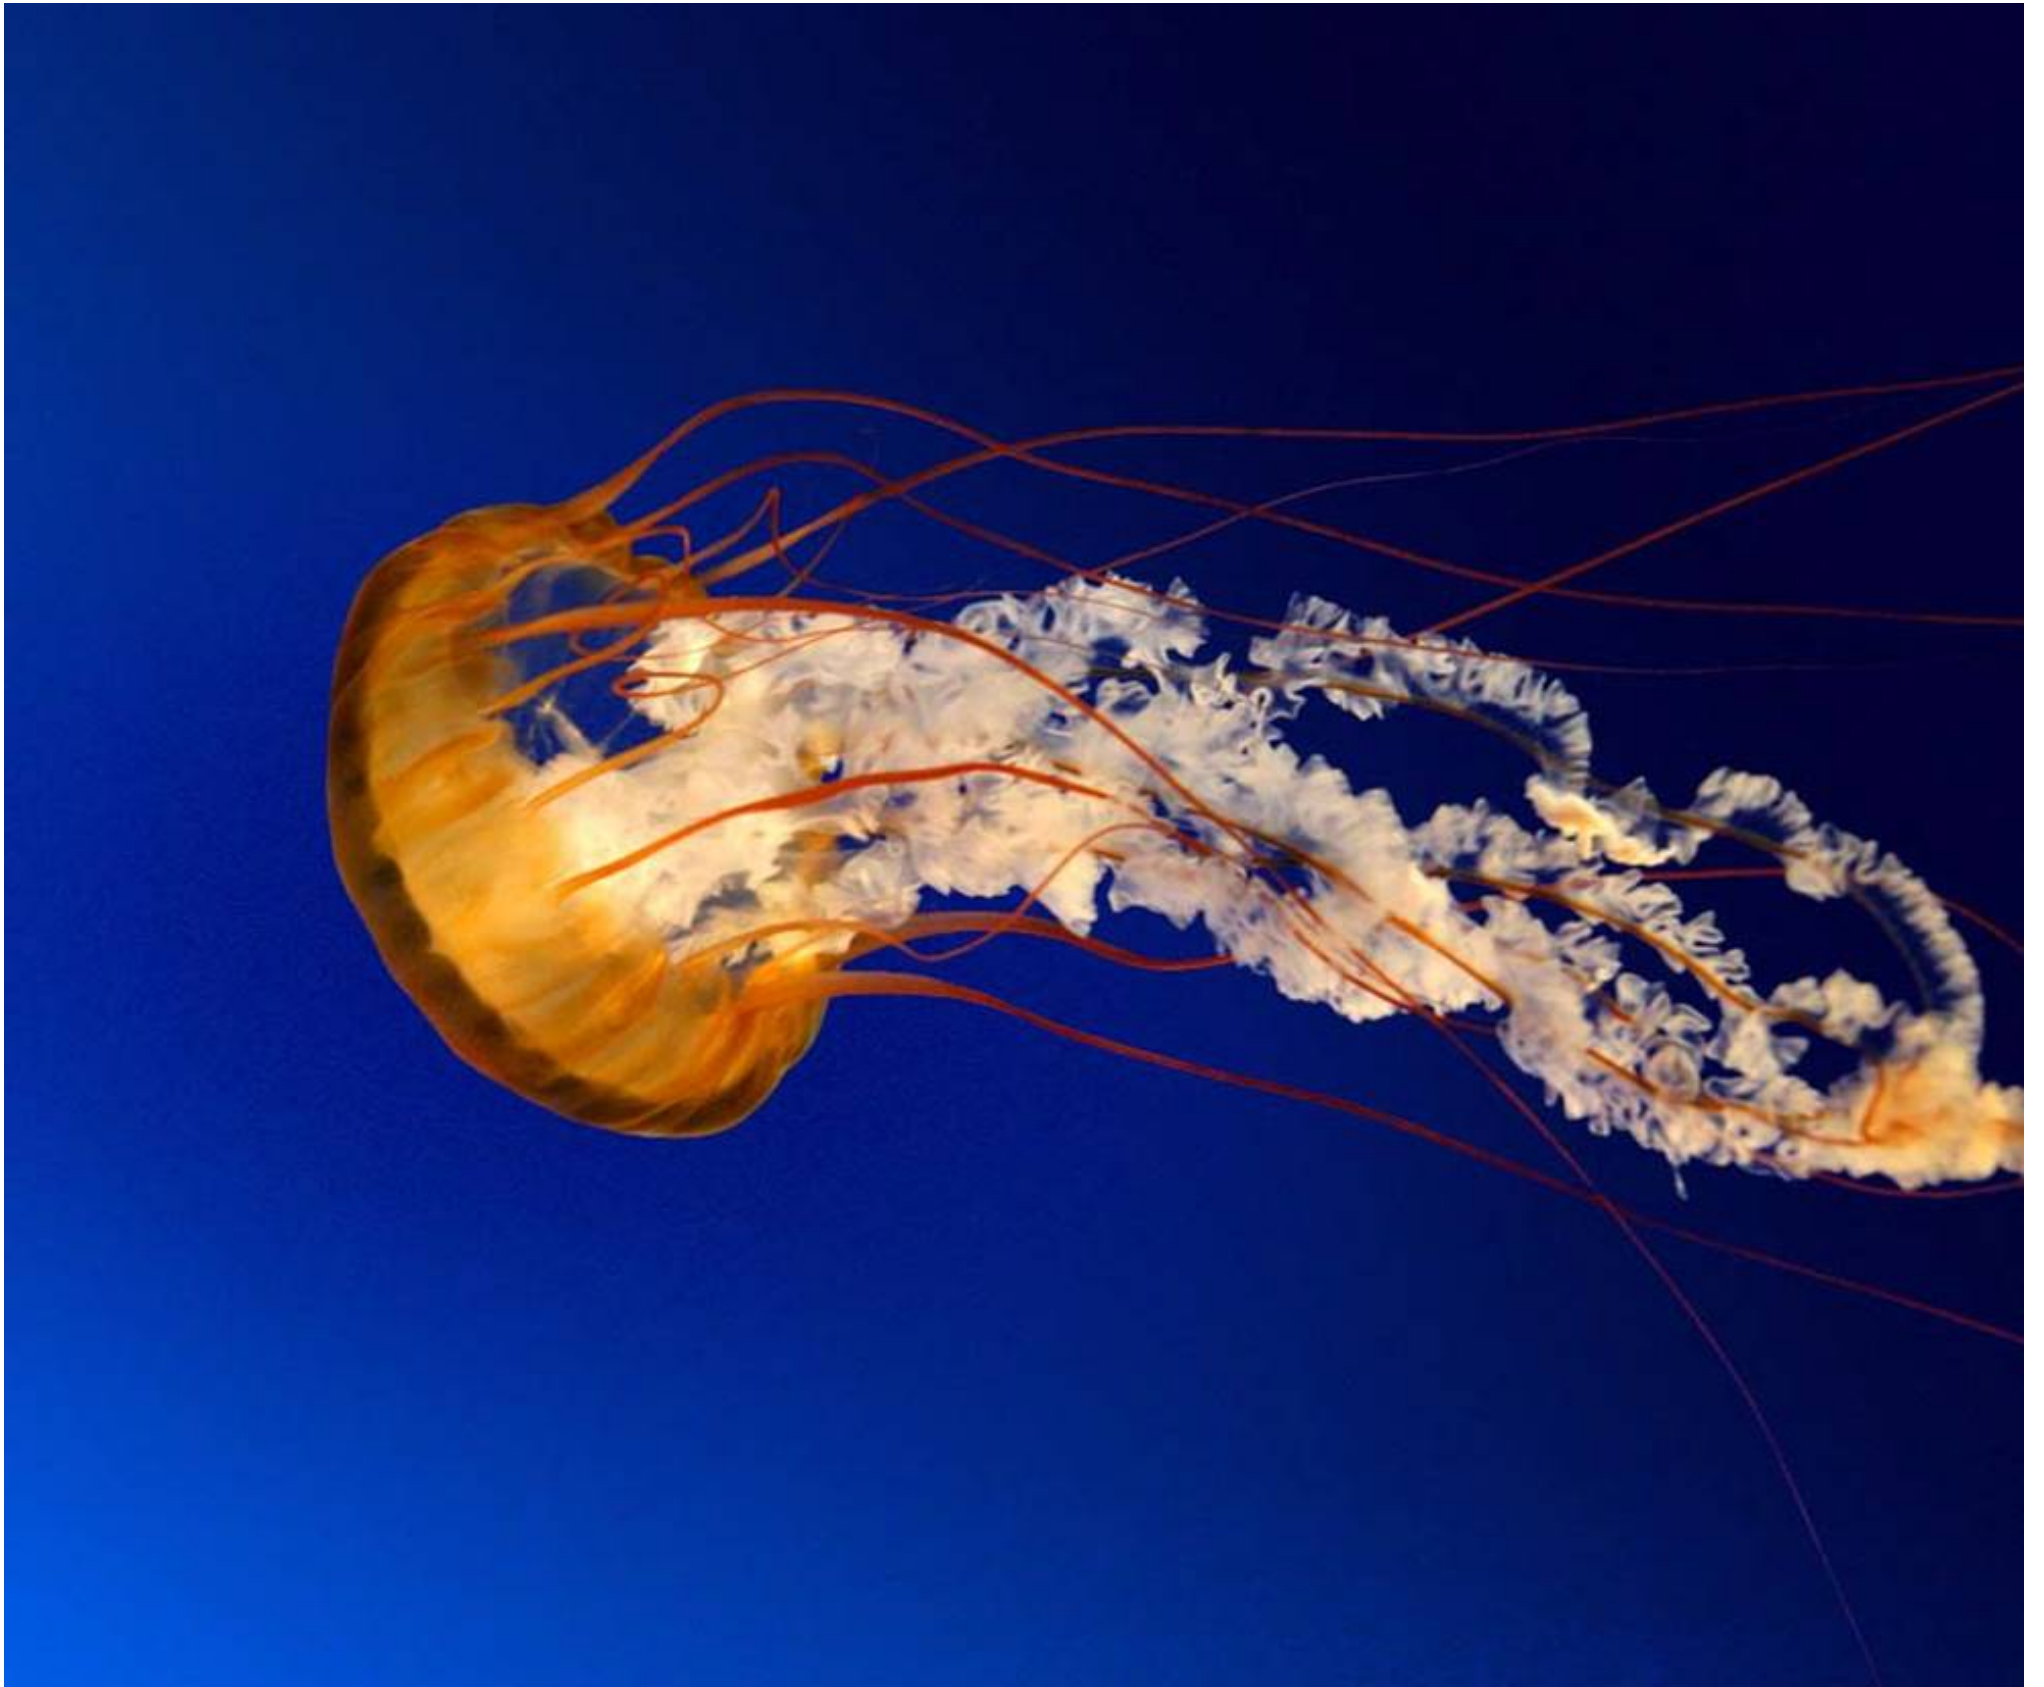
Answer: B

Explanation:

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

NEW QUESTION 123
- (Topic 2)
A user has configured the AWS CloudWatch alarm for estimated usage charges in the US East region. Which of the below mentioned statements is not true with respect to the estimated charges?
Exhibit:

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

Answer: D

Explanation:

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

NEW QUESTION 124
- (Topic 2)
A user is trying to understand AWS SNS. To which of the below mentioned end points is SNS unable to send a notification?

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

Answer: D

Explanation:

Amazon Simple Notification Service (Amazon SNS. is a fast, flexible, and fully managed push messaging service. Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS. queues or to any HTTP endpoint. The user can select one the following transports as part of the subscription requests: "HTTP", "HTTPS","Email", "Email-JSON", "SQS", "and SMS".

NEW QUESTION 126
- (Topic 2)
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. Stop one of the instances and change the availability zone
B. The zone can only be modified using the AWS CLI
C. From the AWS EC2 console, select the Actions - > Change zones and specify new zone
D. Create an AMI of the running instance and launch the instance in a separate AZ

Answer: D

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 129
- (Topic 2)
A user has launched an EBS backed EC2 instance. What will be the difference while performing the restart or stop/start options on that instance?

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

Answer: A

Explanation:

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

NEW QUESTION 134
- (Topic 3)
A user has launched an EC2 instance from an instance store backed AMI. The user has attached an additional instance store volume to the instance. The user wants to create an AMI from the running instance. Will the AMI have the additional instance store volume data?

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

Answer: A

Explanation:

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

NEW QUESTION 138
- (Topic 3)
The CFO of a company wants to allow one of his employees to view only the AWS usage report page. Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

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

Answer: C

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": [
"aws-portal:ViewUsage"
],
"Resource": "*"
}
]
}

NEW QUESTION 143
- (Topic 3)
A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services does not provide detailed monitoring with CloudWatch?

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

Answer: A

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. Services, such as RDS, EC2, Auto Scaling, ELB, and Route 53 can provide the monitoring data every minute.

NEW QUESTION 147
- (Topic 3)
A user is trying to connect to a running EC2 instance using SSH. However, the user gets an Unprotected
Private Key File error. Which of the below mentioned options can be a possible reason for rejection?

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

Answer: A

Explanation:

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

NEW QUESTION 151
- (Topic 3)
A user has created an Auto Scaling group with default configurations from CLI. The user wants to setup the CloudWatch alarm on the EC2 instances, which are launched by the Auto Scaling group. The user has setup an alarm to monitor the CPU utilization every minute. Which of the below mentioned statements is true?

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

Answer: B

Explanation:

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

NEW QUESTION 153
- (Topic 3)
Your mission is to create a lights-out datacenter environment, and you plan to use AWS OpsWorks to accomplish this. First you created a stack and added an App Server layer with an instance running in it. Next you added an application to the instance, and now you need to deploy a MySQL RDS database instance.
Which of the following answers accurately describe how to add a backend database server to an OpsWorks stack? Choose 3 answers

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

Answer: BCE

NEW QUESTION 157
- (Topic 3)
A user has launched an RDS MySQL DB with the Multi AZ feature. The user has scheduled the scaling of

instance storage during maintenance window. What is the correct order of events during maintenance window?
Perform maintenance on standby
Promote standby to primary
Perform maintenance on original primary
Promote original master back as primary

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

Answer: B

Explanation:

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

NEW QUESTION 161
- (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 162
- (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 164
- (Topic 3)
An organization has setup Auto Scaling with ELB. Due to some manual error, one of the instances got
rebooted. Thus, it failed the Auto Scaling health check. Auto Scaling has marked it for replacement. How can the system admin ensure that the instance does not get terminated?

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

Answer: D

Explanation:

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

NEW QUESTION 168
- (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 173
- (Topic 3)
A user has created a VPC with public and private subnets using the VPC wizard. The VPC has CIDR
20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24 . The NAT instance ID is i-a12345. Which of the below mentioned entries are required in the main route table attached with the private subnet to allow instances to connect with the internet?

A. Destination: 0.0.0.0/0 and Target: i-a12345
B. Destination: 20.0.0.0/0 and Target: 80
C. Destination: 20.0.0.0/0 and Target: i-a12345
D. Destination: 20.0.0.0/24 and Target: i-a12345

Answer: A

Explanation:

A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create two route tables and attach to the subnets. The main route table will have the entry "Destination: 0.0.0.0/0 and Target: ia12345", which allows all the instances in the private subnet to connect to the internet using NAT.

NEW QUESTION 174
- (Topic 3)
A user wants to upload a complete folder to AWS S3 using the S3 Management console. How can the user perform this activity?

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

Answer: D

Explanation:

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

NEW QUESTION 177
- (Topic 3)
A user has created a public subnet with VPC and launched an EC2 instance within it. The user is trying to
delete the subnet. What will happen in this scenario?

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

Answer: B

Explanation:

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

NEW QUESTION 179
- (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 183
- (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 185
- (Topic 3)
A sys admin has enabled a log on ELB. Which of the below mentioned activities are not captured by the log?

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

Answer: B

Explanation:

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

NEW QUESTION 188
- (Topic 3)
A user is using Cloudformation to launch an EC2 instance and then configure an application after the instance is launched. The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly. How can the user configure this?

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

Answer: D

Explanation:

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

NEW QUESTION 192
- (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 193
- (Topic 3)
A user has launched an EC2 instance from an instance store backed AMI. If the user restarts the instance, what will happen to the ephermal storage data?

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

Answer: D

Explanation:

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

NEW QUESTION 195
- (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 196
- (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 201
- (Topic 3)
When you put objects in Amazon S3, what is the indication that an object was successfully stored?

A. Each S3 account has a special bucket named_s3_log

B. Success codes are written to this bucket with a timestamp and checksu
C. A success code is inserted into the S3 object metadat
D. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successfu
E. Amazon S3 is engineered for 99.999999999% durabilit
F. Therefore there is no need to confirm that data was inserte

Answer: B

Explanation:
 Reference:
http://docs.aws.amazon.com/AmazonS3/latest/API/RESTObjectPUT.html

NEW QUESTION 206
- (Topic 3)
A user is trying to create an EBS volume with the highest PIOPS supported by EBS. What is the minimum size of EBS required to have the maximum IOPS?

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

Answer: C

Explanation:

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

NEW QUESTION 208
- (Topic 3)
An organization has launched 5 instances: 2 for production and 3 for testing. The organization wants that one particular group of IAM users should only access the test instances and not the production ones. How can the organization set that as a part of the policy?

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

Answer: D

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The user can add conditions as a part of the IAM policies. The condition can be set on AWS Tags, Time, and Client IP as well as on various parameters. If the organization wants the user to access only specific instances he should define proper tags and add to the IAM policy condition.
The sample policy is shown below.
"Statement": [
{
"Action": "ec2:*",
"Effect": "Allow",
"Resource": "*",
"Condition": {
"StringEquals": {
"ec2:ResourceTag/InstanceType": "Production"
}
}
}
]

NEW QUESTION 213
- (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 217

- (Topic 3)
A user has configured an HTTPS listener on an ELB. The user has not configured any security policy which can help to negotiate SSL between the client and ELB. What will ELB do in this scenario?

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

Answer: B

Explanation:

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

NEW QUESTION 221
- (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 225
- (Topic 3)
A user is measuring the CPU utilization of a private data centre machine every minute. The machine provides the aggregate of data every hour, such as Sum of data", "Min value", "Max value, and "Number of Data points".
The user wants to send these values to CloudWatch. How can the user achieve this?

A. Send the data using the put-metric-data command with the aggregate-values parameter
B. Send the data using the put-metric-data command with the average-values parameter
C. Send the data using the put-metric-data command with the statistic-values parameter
D. Send the data using the put-metric-data command with the aggregate –data parameter

Answer: C

Explanation:

AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. The user can publish the data to CloudWatch as single data points or as an aggregated set of data points called a statistic set using the command put-metric-data. When sending the aggregate data, the user needs to send it with the parameter statistic-values: awscloudwatch put-metric-data --metric-name <Name> --namespace <Custom namespace> --timestamp <UTC Format> --statistic-values Sum=XX,Minimum=YY,Maximum=AA,SampleCount=BB --unit Milliseconds

NEW QUESTION 230
- (Topic 3)
A user has created a VPC with public and private subnets. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.1.0/24 and the public subnet uses CIDR 20.0.0.0/24. The user is planning to host a web server in the public subnet (port 80. and a DB server in the private subnet (port 3306.. The user is configuring a security group of the NAT instance. Which of the below mentioned entries is not required for the NAT security group?

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

Answer: C

Explanation:

A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet to host the web server and DB server respectively, the user should configure that the instances in the private subnet can connect to the internet using the NAT instances. The user should first configure that NAT can receive traffic on ports 80 and 443 from the private subnet. Thus, allow ports 80 and 443 in Inbound for the private subnet 20.0.1.0/24. Now to route this traffic to the internet configure ports 80 and Amazon AWS-SysOps : Practice Test
443 in Outbound with destination 0.0.0.0/0. The NAT should not have an entry for the public subnet CIDR.

NEW QUESTION 235
- (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 236
- (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 240
- (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 244
- (Topic 3)
A user is trying to create a PIOPS EBS volume with 4000 IOPS and 100 GB size. AWS does not allow the user to create this volume. What is the possible root cause for this?

A. The ratio between IOPS and the EBS volume is higher than 30
B. The maximum IOPS supported by EBS is 3000
C. The ratio between IOPS and the EBS volume is lower than 50
D. PIOPS is supported for EBS higher than 500 GB size

Answer: A

Explanation:

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

NEW QUESTION 246
- (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 250
- (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 254
- (Topic 3)
A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered. What time out period should the user specify for connection draining?

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

Answer: B

NEW QUESTION 258
- (Topic 3)
A user has configured ELB with a TCP listener at ELB as well as on the back-end instances. The user wants to enable a proxy protocol to capture the source and destination IP information in the header. Which of the below mentioned statements helps the user understand a proxy protocol with TCP configuration?

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

Answer: A

Explanation:

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

NEW QUESTION 262
- (Topic 3)
A user has configured an SSL listener at ELB as well as on the back-end instances. Which of the below mentioned statements helps the user understand ELB traffic handling with respect to the SSL listener?

A. It is not possible to have the SSL listener both at ELB and back-end instances
B. ELB will modify headers to add requestor details
C. ELB will intercept the request to add the cookie details if sticky session is enabled

D. ELB will not modify the headers

Answer: D

Explanation:

When the user has configured Transmission Control Protocol (TCP. or Secure Sockets Layer (SSL. for both front-end and back-end connections of the Elastic Load Balancer, the load balancer forwards the request to the back-end instances without modifying the request headers unless the proxy header is enabled. SSL does not support sticky sessions. If the user has enabled a proxy protocol it adds the source and destination IP to the header.

NEW QUESTION 265
- (Topic 3)
In order to optimize performance for a compute cluster that requires low inter-node latency, which feature in the following list should you use?

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

Answer: D

NEW QUESTION 266
- (Topic 3)
A user is using a small MySQL RDS DB. The user is experiencing high latency due to the Multi AZ feature.Which of the below mentioned options may not help the user in this situation?

A. Schedule the automated back up in non-working hours
B. Use a large or higher size instance
C. Use PIOPS
D. Take a snapshot from standby Replica

Answer: D

Explanation:

An RDS DB instance which has enabled Multi AZ deployments may experience increased write and commit latency compared to a Single AZ deployment, due to synchronous data replication. The user may also face changes in latency if deployment fails over to the standby replica. For production workloads, AWS recommends the user to use provisioned IOPS and DB instance classes (m1.large and larger. as they are optimized for provisioned IOPS to give a fast, and consistent performance. With Multi AZ feature, the user can not have option to take snapshot from replica.

NEW QUESTION 269
- (Topic 3)
A storage admin wants to encrypt all the objects stored in S3 using server side encryption. The user does not want to use the AES 256 encryption key provided by S3. How can the user achieve this?

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

Answer: D

Explanation:

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

NEW QUESTION 271
- (Topic 3)
A .NET application that you manage is running in Elastic Beanstalk. Your developers tell you they will need access to application log files to debug issues that arise. The infrastructure will scale up and down.
How can you ensure the developers will be able to access only the log files?

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

Answer: D

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

NEW QUESTION 275

- (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 279
- (Topic 3)
A user has created a VPC with a subnet and a security group. The user has launched an instance in that
subnet and attached a public IP. The user is still unable to connect to the instance. The internet gateway has also been created. What can be the reason for the error?

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

Answer: A

Explanation:

A Virtual Private Cloud (VPC. is a virtual network dedicated to the user's AWS account. AWS provides two features the user can use to increase security in VPC: security groups and network ACLs. Security groups work at the instance level. When a user launches an instance and wants to connect to an instance, he needs an internet gateway. The internet gateway should be configured with the route table to allow traffic from the internet.

NEW QUESTION 283
- (Topic 3)
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. The user's data centre has CIDR 172.28.0.0/12. The user has also setup a NAT instance (i-123456. to allow traffic to the internet from the VPN subnet. Which of the below mentioned options is not a valid entry for the main route table in this scenario?

A. Destination: 20.0.1.0/24 and Target: i-12345
B. Destination: 0.0.0.0/0 and Target: i-12345
C. Destination: 172.28.0.0/12 and Target: vgw-12345
D. Destination: 20.0.0.0/16 and Target: local

Answer: A

Explanation:

The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, 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. If the user has setup a NAT instance to route all the internet requests then all requests to the internet should be routed to it. All requests to the organization's DC will be routed to the VPN gateway. Here are the valid entries for the main route table in this scenario: Destination: 0.0.0.0/0 & Target: i-12345 (To route all internet traffic to the NAT Instance. Destination: 172.28.0.0/12 & Target: vgw-12345 (To route all the organization's data centre traffic to the VPN gateway. Destination: 20.0.0.0/16 & Target: local (To allow local routing in VPC.

NEW QUESTION 286
- (Topic 3)
An organization is measuring the latency of an application every minute and storing data inside a file in the JSON format. The organization wants to send all latency data to AWS CloudWatch. How can the organization achieve this?

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

Answer: C

Explanation:

AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. The user has to always include the namespace as part of the request. If the user wants to upload the custom data from a Amazon AWS-SysOps : Practice Test file, he can supply file name along with the parameter -- metric-data to command put-metric-data.

NEW QUESTION 290
- (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 293
- (Topic 3)
A user has created a VPC with CIDR 20.0.0.0/16 using VPC Wizard. The user has created a public CIDR
(20.0.0.0/24. and a VPN only subnet CIDR (20.0.1.0/24. along with the hardware VPN access to connect to the user's data centre. Which of the below mentioned components is not present when the VPC is setup with the wizard?

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

Answer: B

Explanation:

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

NEW QUESTION 297
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-SysOps Practice Exam Features:

* AWS-SysOps Questions and Answers Updated Frequently

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

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

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

100% Actual & Verified — Instant Download, Please Click
Order The AWS-SysOps Practice Test Here