AWS-SysOps Dumps

Amazon AWS Certified SysOps Administrator - Associate

https://www.certleader.com/AWS-SysOps-dumps.html
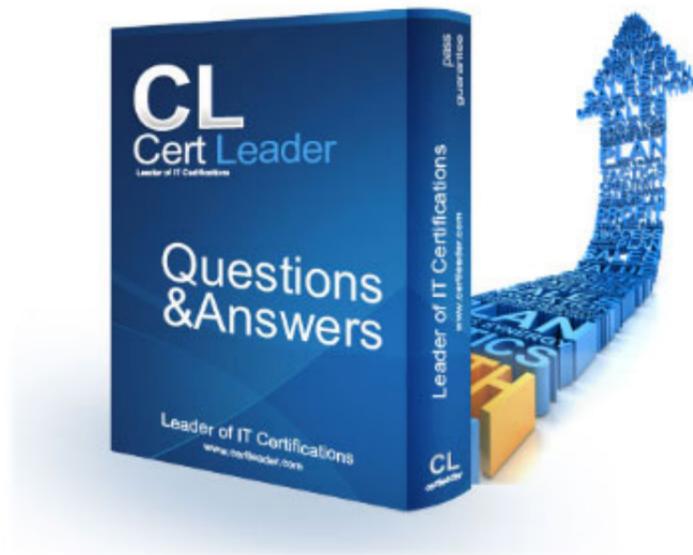
NEW QUESTION 1
- (Topic 1)
A media company produces new video files on-premises every day with a total size of around 100GBS after compression All files have a size of 1 -2 GB and need to be uploaded to Amazon S3 every night in a fixed time window between 3am and 5am Current upload takes almost 3 hours, 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 2
- (Topic 1)
How can the domain's zone apex for example "myzoneapexdomain com" be pointed towards an Elastic Load Balancer?

A. By using an AAAA record
B. By using an A record
C. By using an Amazon Route 53 CNAME record
D. By using an Amazon Route 53 Alias record

Answer: D

Explanation:
Reference:
http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html

NEW QUESTION 3
- (Topic 1)
You have been asked to leverage Amazon VPC BC2 and SOS to implement an application that submits and receives millions of messages per second to a message queue. You want to ensure your application has sufficient bandwidth between your EC2 instances and SQS Which option will provide the most scalable solution for communicating between the application and SQS?

A. Ensure the application instances are properly configured with an Elastic Load Balancer
B. Ensure the application instances are launched in private subnets with the EBS-optimized option enabled
C. Ensure the application instances are launched in public subnets with the associate-public-IP-address=true option enabled
D. Launch application instances in private subnets with an Auto Scaling group and Auto Scaling triggers configured to watch the SQS queue size

Answer: B

Explanation:
Reference:
http://www.cardinalpath.com/autoscaling-your-website-with-amazon-web-services-part-2/

NEW QUESTION 4
- (Topic 1)
You have a web-style application with a stateless but CPU and memory-intensive web tier running on a cc2 8xlarge EC2 instance inside of a VPC The instance when under load is having problems returning requests within the SLA as defined by your business The application maintains its state in a DynamoDB table, but the data tier is properly provisioned and responses are consistently fast.
How can you best resolve the issue of the application responses not meeting your SLA?

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

Answer: B

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

NEW QUESTION 5
- (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 6
- (Topic 1)
You have decided to change the Instance type for instances running In your application tier that are using Auto Scaling.
In which area below would you change the instance type definition?

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

Answer: A

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

NEW QUESTION 7
- (Topic 1)
You are designing a system that has a Bastion host. This component needs to be highly available without human intervention.
Which of the following approaches would you select?

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

Answer: C

NEW QUESTION 8
- (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 9
- (Topic 1)
When assessing an organization s use of AWS API access credentials which of the following three credentials should be evaluated?
Choose 3 answers

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

Answer: ACD

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

NEW QUESTION 10
- (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 have a web application leveraging an Elastic Load Balancer (ELB) In front of the web servers deployed using an Auto Scaling Group Your database is running on Relational
Database Service (RDS) The application serves out technical articles and responses to them in general there are more views of an article than there are responses to the article. On occasion, an article on the site becomes extremely popular resulting in significant traffic Increases that causes the site to go down.
What could you do to help alleviate the pressure on the infrastructure while maintaining availability during these events?
Choose 3 answers

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

Answer: ACE

NEW QUESTION 15
- (Topic 1)
What is a placement group?

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

Answer: B

Explanation:
 Reference:
http://aws.amazon.com/ec2/faqs/

NEW QUESTION 20
- (Topic 1)
The majority of your Infrastructure is on premises and you have a small footprint on AWS Your company has decided to roll out a new application that is heavily dependent on low latency connectivity to LOAP for authentication Your security policy requires minimal changes to the company's existing application user management processes.
What option would you implement to successfully launch this application1?

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

Answer: D

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

NEW QUESTION 21
- (Topic 1)
Your application currently leverages AWS Auto Scaling to grow and shrink as load Increases/ decreases and has been performing well Your marketing team expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks Your forecast for the approximate number of Amazon EC2 instances necessary to meet the peak demand is 175.
What should you do to avoid potential service disruptions during the ramp up in traffic?

A. Ensure that you have pre-allocated 175 Elastic IP addresses so that each server will be able to obtain one as it launches
B. Check the service limits in Trusted Advisor and adjust as necessary so the forecasted count remains within limit
C. Change your Auto Scaling configuration to set a desired capacity of 175 prior to the launch of the marketing campaign
D. Pre-warm your Elastic Load Balancer to match the requests per second anticipated during peak demand prior to the marketing campaign

Answer: D

NEW QUESTION 24
- (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 25
- (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 27
- (Topic 1)
If you want to launch Amazon Elastic Compute Cloud (EC2) Instances and assign each Instance a predetermined private IP address you should:

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

Answer: C

Explanation:
 Reference:
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-ip-addressing.html

NEW QUESTION 31
- (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 36
- (Topic 1)
Your company Is moving towards tracking web page users with a small tracking
Image loaded on each page Currently you are serving this image out of US-East, but are starting to get concerned about the time It takes to load the image for users on the west coast.
What are the two best ways to speed up serving this image?
Choose 2 answers

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

Answer: AD

NEW QUESTION 41
- (Topic 1)
You have started a new job and are reviewing your company's infrastructure on AWS You notice one web application where they have an Elastic Load Balancer (&B) in front of web instances in an Auto Scaling Group When you check the metrics for the ELB in CloudWatch you see four healthy instances In Availability Zone (AZ) A and zero in AZ B There are zero unhealthy instances.
What do you need to fix to balance the instances across AZs?

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

Answer: B

NEW QUESTION 44
- (Topic 1)
Which two AWS services provide out-of-the-box user configurable automatic backup-as-a-
service and backup rotation options?
Choose 2 answers

A. Amazon S3
B. Amazon RDS
C. Amazon EBS
D. Amazon Red shift

Answer: BD

NEW QUESTION 45
- (Topic 1)
You are using ElastiCache Memcached to store session state and cache database queries in your infrastructure. You notice in CloudWatch that Evictions and GetMisses are Doth very high.
What two actions could you take to rectify this?
Choose 2 answers

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

Answer: BD

NEW QUESTION 49
- (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 53
- (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 55
- (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 60
- (Topic 1)
When creation of an EBS snapshot Is initiated but not completed the EBS volume?

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

Answer: C

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

NEW QUESTION 65
- (Topic 2)
A user is planning to setup infrastructure on AWS for the Christmas sales. The user is planning to use Auto Scaling based on the schedule for proactive scaling.
What advise would you give to the user?

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

Answer: C

Explanation:

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

NEW QUESTION 70
- (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 74
- (Topic 2)
A user has configured an Auto Scaling group with ELB. The user has enabled detailed CloudWatch monitoring on Auto Scaling. Which of the below mentioned statements will help the user understand the functionality better?

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

Answer: B

Explanation:

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. Auto Scaling includes 7 metrics and 1 dimension, and sends data to CloudWatch every 5 minutes by default. The user can enable detailed monitoring for Auto Scaling, which sends data to CloudWatch every minute. However, this will have some extra-costs.

NEW QUESTION 78
- (Topic 2)
An organization is generating digital policy files which are required by the admins for verification. Once the files are verified they may not be required in the future unless there is some compliance issue. If the organization wants to save them in a cost effective way, which is the best possible solution?

A. AWS RRS

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

Answer: D

Explanation:

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

NEW QUESTION 79
- (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 a versa. The bucket policy can be defined at the bucket level which allows the objects as well as the bucket to be public with a single policy applied to that bucket. It cannot be applied at the object level. The folders are similar to objects with no content. Thus, folders can have only ACL and cannot have a policy.

NEW QUESTION 82
- (Topic 2)
A user has created an ELB with the availability zone US-East-1A. The user wants to add more zones to ELB to achieve High Availability. How can the user add more zones to the existing ELB?

A. It is not possible to add more zones to the existing ELB
B. The only option is to launch instances in different zones and add to ELB
C. The user should stop the ELB and add zones and instances as required
D. The user can add zones on the fly from the AWS console

Answer: D

Explanation:

The user has created an Elastic Load Balancer with the availability zone and wants to add more zones to the existing ELB. The user can do so in two ways:
From the console or CLI, add new zones to ELB;
Launch instances in a separate AZ and add instances to the existing ELB.

NEW QUESTION 84
- (Topic 2)
A user has received a message from the support team that an issue occurred 1 week back between 3 AM to 4 AM and the EC2 server was not reachable. The user is checking the CloudWatch metrics of that instance. How can the user find the data easily using the CloudWatch console?

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

Answer: D

Explanation:

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

NEW QUESTION 89
- (Topic 2)
An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

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

Answer: C

Explanation:

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

NEW QUESTION 90
- (Topic 2)
A user has created an ELB with three instances. How many security groups will ELB create by default?

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

Answer: C

Explanation:

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

NEW QUESTION 91
- (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 94
- (Topic 2)
A user has setup a billing alarm using CloudWatch for $200. The usage of AWS exceeded $200 after some days. The user wants to increase the limit from $200 to $400? What should the user do?

A. Create a new alarm of $400 and link it with the first alarm
B. It is not possible to modify the alarm once it has crossed the usage limit
C. Update the alarm to set the limit at $400 instead of $200
D. Create a new alarm for the additional $200 amount

Answer: C

Explanation:

AWS CloudWatch supports enabling the billing alarm on the total AWS charges. 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. This data also includes the estimated charges for every service in AWS used by the user, as well as the estimated overall AWS charges. If the user wants to increase the limit, the user can modify the alarm and specify a new threshold.

NEW QUESTION 99
- (Topic 2)
A user has launched an EC2 instance from an instance store backed AMI. The infrastructure team wants to create an AMI from the running instance. Which of the below mentioned steps will not be performed while creating the AMI?

A. Define the AMI launch permissions
B. Upload the bundled volume
C. Register the AMI
D. Bundle the volume

Answer: A

Explanation:

When the user has launched an EC2 instance from an instance store backed AMI, it will need to follow certain steps, such as "Bundling the root volume", "Uploading the bundled volume" and "Register the AMI". Once the AMI is created the user can setup the launch permission. However, it is not required to setup during the launch.

NEW QUESTION 102
- (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 103
- (Topic 2)
An organization is planning to use AWS for 5 different departments. The finance department is responsible to pay for all the accounts. However, they want the cost separation for each account to map with the right cost centre. How can the finance department achieve this?

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

Answer: A

Explanation:

AWS consolidated billing enables the organization to consolidate payments for multiple Amazon Web Services (AWS. accounts within a single organization by making a single paying account. Consolidated billing enables the organization to see a combined view of the AWS charges incurred by each account as well as obtain a detailed cost report for each of the individual AWS accounts associated with the paying account.

NEW QUESTION 107
- (Topic 2)
An organization is planning to use AWS for their production roll out. The organization wants to implement
automation for deployment such that it will automatically create a LAMP stack, download the latest PHP
installable from S3 and setup the ELB. Which of the below mentioned AWS services meets the quirement for making an orderly deployment of the software?

A. AWS Elastic Beanstalk
B. AWS Cloudfront
C. AWS Cloudformation
D. AWS DevOps

Answer: C

Explanation:

AWS Cloudformation is an application management tool which provides application modelling, deployment, configuration, management and related activities. Cloudformation provides an easy way to create and delete the collection of related AWS resources and provision them in an orderly way. AWS CloudFormation automates and simplifies the task of repeatedly and predictably creating groups of related resources that power the user's applications. AWS Cloudfront is a CDN; Elastic Beanstalk does quite a few of the required tasks. However, it is a PAAS which uses a ready AMI. AWS Elastic Beanstalk provides an environment to easily develop and run applications in the cloud.

NEW QUESTION 112
- (Topic 2)
A user has created a VPC with CIDR 20.0.0.0/24. The user has created a public subnet with CIDR 20.0.0.0/25. The user is trying to create the private subnet with CIDR 20.0.0.128/25. Which of the below mentioned statements is true in this scenario?

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

Answer: B

Explanation:

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

NEW QUESTION 113
- (Topic 2)
A user has configured Elastic Load Balancing by enabling a Secure Socket Layer (SSL. negotiation
configuration known as a Security Policy. Which of the below mentioned options is not part of this secure policy while negotiating the SSL connection between the user and the client?

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

Answer: B

Explanation:

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

NEW QUESTION 116
- (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 119
- (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 121
- (Topic 2)
A user has created an S3 bucket which is not publicly accessible. The bucket is having thirty objects which are also private. If the user wants to make the objects public, how can he configure this with minimal efforts?

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

Answer: C

Explanation:

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

NEW QUESTION 126
- (Topic 2)
A user is checking the CloudWatch metrics from the AWS console. The user notices that the CloudWatch data is coming in UTC. The user wants to convert the data to a local time zone. How can the user perform this?

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

Answer: B

Explanation:

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

NEW QUESTION 130
- (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 133
- (Topic 2)
A user has created a web application with Auto Scaling. The user is regularly monitoring the application and he observed that the traffic is highest on Thursday and Friday between 8 AM to 6 PM. What is the best solution to handle scaling in this case?

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

Answer: B

Explanation:

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

NEW QUESTION 135
- (Topic 2)
A user is trying to save some cost on the AWS services. Which of the below mentioned options will not help him save cost?

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

Answer: B

Explanation:

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

NEW QUESTION 140
- (Topic 2)
A user is trying to configure the CloudWatch billing alarm. Which of the below mentioned steps should be
performed by the user for the first time alarm creation in the AWS Account Management section?

A. Enable Receiving Billing Reports
B. Enable Receiving Billing Alerts
C. Enable AWS billing utility
D. Enable CloudWatch Billing Threshold

Answer: B

Explanation:

AWS CloudWatch supports enabling the billing alarm on the total AWS charges. Before the user can create an alarm on the estimated charges, he must enable monitoring of the estimated AWS charges, by selecting the option "Enable receiving billing alerts". It takes about 15 minutes before the user can view the billing data. The user can then create the alarms.

NEW QUESTION 145
- (Topic 2)
A sys admin is trying to understand the Auto Scaling activities. Which of the below mentioned processes is not performed by Auto Scaling?

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

Answer: A

Explanation:

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

NEW QUESTION 149
- (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 152
- (Topic 2)
A user has setup a web application on EC2. The user is generating a log of the application performance at every second. There are multiple entries for each second. If the user wants to send that data to CloudWatch every minute, what should he do?

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

Answer: C

Explanation:

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

NEW QUESTION 155
- (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 158
- (Topic 2)
A user has setup Auto Scaling with ELB on the EC2 instances. The user wants to configure that whenever the CPU utilization is below 10%, Auto Scaling should remove one instance. How can the user configure this?

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

Answer: D

Explanation:

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

NEW QUESTION 161
- (Topic 2)
A sys admin has created a shopping cart application and hosted it on EC2. The EC2 instances are running behind ELB. The admin wants to ensure that the end user request will always go to the EC2 instance where the user session has been created. How can the admin configure this?

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

Answer: C

Explanation:

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

NEW QUESTION 162
- (Topic 2)
A user is trying to delete an Auto Scaling group from CLI. Which of the below mentioned steps are to be performed by the user?

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

Answer: C

Explanation:

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

NEW QUESTION 167
- (Topic 2)
A user has a refrigerator plant. The user is measuring the temperature of the plant every 15 minutes. If the user wants to send the data to CloudWatch to view the data visually, which of the below mentioned statements is true with respect to the information given above?

A. The user needs to use AWS CLI or API to upload the data
B. The user can use the AWS Import Export facility to import data to CloudWatch
C. The user will upload data from the AWS console
D. The user cannot upload data to CloudWatch since it is not an AWS service metric

Answer: A

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. While sending the data the user has to include the metric name, namespace and timezone as part of the request.

NEW QUESTION 172

- (Topic 2)
A user is publishing custom metrics to CloudWatch. Which of the below mentioned statements will help the user understand the functionality better?

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

Answer: B

Explanation:

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

NEW QUESTION 176
- (Topic 2)
A user has configured a VPC with a new subnet. The user has created a security group. The user wants to configure that instances of the same subnet communicate with each other. How can the user configure this with the security group?

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

Answer: C

Explanation:

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

NEW QUESTION 180
- (Topic 2)
A system admin is planning to setup event notifications on RDS. Which of the below mentioned services will help the admin setup notifications?

A. AWS SES
B. AWS Cloudtrail
C. AWS Cloudwatch
D. AWS SNS

Answer: D

Explanation:

Amazon RDS uses the Amazon Simple Notification Service to provide a notification when an Amazon RDS event occurs. These notifications can be in any notification form supported by Amazon SNS for an AWS region, such as an email, a text message or a call to an HTTP endpoint

NEW QUESTION 182
- (Topic 2)
A user has created a VPC with CIDR 20.0.0.0/24. The user has created a public subnet with CIDR 20.0.0.0/25 and a private subnet with CIDR 20.0.0.128/25. The user has launched one instance each in the private and public subnets. Which of the below mentioned options cannot be the correct IP address (private IP. assigned to an instance in the public or private subnet?

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

Answer: A

Explanation:

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

NEW QUESTION 185
- (Topic 3)
A user has deployed an application on an EBS backed EC2 instance. For a better performance of application, it requires dedicated EC2 to EBS traffic. How can the user achieve this?

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

D. Launch the EC2 instance as EBS optimized with PIOPS EBS

Answer: D

Explanation:

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

NEW QUESTION 187
- (Topic 3)
A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling terminate process only for a while. What will happen to the availability zone rebalancing process (AZRebalance. during this period?

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

Answer: B

Explanation:

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

NEW QUESTION 188
- (Topic 3)
A user is trying to understand the CloudWatch metrics for the AWS services. It is required that the user should first understand the namespace for the AWS services. Which of the below mentioned is not a valid namespace for the AWS services?

A. AWS/StorageGateway
B. AWS/CloudTrail
C. AWS/ElastiCache
D. AWS/SWF

Answer: B

Explanation:

Amazon CloudWatch is basically a metrics repository. The AWS product puts metrics into this repository, and the user can retrieve the data or statistics based on those metrics. To distinguish the data for each service, the CloudWatch metric has a namespace. Namespaces are containers for metrics. All AWS services that provide the Amazon CloudWatch data use a namespace string, beginning with "AWS/". All the services which are supported by CloudWatch will have some namespace. CloudWatch does not monitor CloudTrail. Thus, the namespace "AWS/CloudTrail" is incorrect.

NEW QUESTION 192
- (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 195

- (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 200
- (Topic 3)
A user has created a VPC with CIDR 20.0.0.0/24. The user has used all the IPs of CIDR and wants to increase the size of the VPC. The user has two subnets:
public (20.0.0.0/28. and private (20.0.1.0/28.. How can the user change the size of the VPC?

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

Answer: B

Explanation:

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

NEW QUESTION 205
- (Topic 3)
A user has launched 5 instances in EC2-CLASSIC and attached 5 elastic IPs to the five different instances in the US East region. The user is creating a VPC in the same region. The user wants to assign an elastic IP to the VPC instance. How can the user achieve this?

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

Answer: D

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

NEW QUESTION 208
- (Topic 3)
A user has setup a VPC with CIDR 20.0.0.0/16. The VPC has a private subnet (20.0.1.0/24. and a public
subnet (20.0.0.0/24.. The user's data centre has CIDR of 20.0.54.0/24 and 20.1.0.0/24. If the private subnet wants to communicate with the data centre, what will happen?

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

Answer: D

Explanation:

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

NEW QUESTION 209
- (Topic 3)
A system admin wants to add more zones to the existing ELB. The system admin wants to perform this activity from CLI. Which of the below mentioned command helps the system admin to add new zones to the existing ELB?

A. elb-enable-zones-for-lb
B. elb-add-zones-for-lb

C. It is not possible to add more zones to the existing ELB
D. elb-configure-zones-for-lb

Answer: A

Explanation:

The user has created an Elastic Load Balancer with the availability zone and wants to add more zones to the existing ELB. The user can do so in two ways:
From the console or CLI, add new zones to ELB;

NEW QUESTION 214
- (Topic 3)
Which of the following statements about this S3 bucket policy is true?

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

A. Denies the server with the IP address 192.166 100.0 full access to the "mybucket" bucket
B. Denies the server with the IP address 192.166 100.188 full access to the "mybucket bucket
C. Grants all the servers within the 192 168 100 0/24 subnet full access to the "mybucket" bucket
D. Grants all the servers within the 192 168 100 188/32 subnet full access to the "mybucket" bucket

Answer: C

NEW QUESTION 219
- (Topic 3)
You have a business-to-business web application running in a VPC consisting of an Elastic Load Balancer (ELB), web servers, application servers and a database.
Your web application should only accept traffic from pre-defined customer IP addresses.
Which two options meet this security requirement? Choose 2 answers A. Configure web server VPC security groups to allow traffic from your customers' IPs

A. Configure your web servers to filter traffic based on the ELB's "X-forwarded-for" header
B. Configure ELB security groups to allow traffic from your customers' IPs and deny all outbound traffic
C. Configure a VPC NACL to allow web traffic from your customers' IPs and deny all outbound traffic

Answer: AB

NEW QUESTION 223
- (Topic 3)
A user has created a VPC with public and private subnets using the VPC wizard. Which of the below
mentioned statements is true in this scenario?

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

Answer: B

Explanation:

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

NEW QUESTION 227
- (Topic 3)
Which method can be used to prevent an IP address block from accessing public objects in an S3 bucket?

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

Answer: A

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

NEW QUESTION 229
- (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 231
- (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 234
- (Topic 3)
A user has enabled detailed CloudWatch monitoring with the AWS Simple Notification Service. Which of the below mentioned statements helps the user understand detailed monitoring better?

A. SNS will send data every minute after configuration
B. There is no need to enable since SNS provides data every minute
C. AWS CloudWatch does not support monitoring for SNS
D. SNS cannot provide data every minute

Answer: D

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. The AWS SNS service sends data every 5 minutes. Thus, it supports only the basic monitoring. The user cannot enable detailed monitoring with SNS.

NEW QUESTION 239
- (Topic 3)
A user has created a VPC with the public and private subnets using the VPC wizard. The VPC has CIDR
20.0.0.0/16. The public subnet uses CIDR 20.0.1.0/24. The user is planning to host a web server in the public subnet (port 80. and a DB server in the private subnet (port 3306.. The user is configuring a security group for the public subnet (WebSecGrp. and the private subnet (DBSecGrp.. Which of the below mentioned entries is required in the private subnet database security group (DBSecGrp.?

A. Allow Inbound on port 3306 for Source Web Server Security Group (WebSecGr
B. Allow Inbound on port 3306 from source 20.0.0.0/16
C. Allow Outbound on port 3306 for Destination Web Server Security Group (WebSecGr
D. Allow Outbound on port 80 for Destination NAT Instance IP

Answer: A

Explanation:

A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet to host the web server and DB server respectively, the user should configure that the instances in the private subnet can receive inbound traffic from the public subnet on the DB port. Thus, configure port 3306 in Inbound with the source as the Web Server Security Group (WebSecGrp.. The user should configure ports 80 and 443 for Destination 0.0.0.0/0 as the route table directs traffic to the NAT instance from the private subnet.

NEW QUESTION 240
- (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 243
- (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 248
- (Topic 3)
An organization has configured Auto Scaling with ELB. One of the instance health check returns the status as Impaired to Auto Scaling. What will Auto Scaling do in this scenario?

A. Perform a health check until cool down before declaring that the instance has failed
B. Terminate the instance and launch a new instance
C. Notify the user using SNS for the failed state
D. Notify ELB to stop sending traffic to the impaired instance

Answer: B

Explanation:

The Auto Scaling group determines the health state of each instance periodically by checking the results of the Amazon EC2 instance status checks. If the instance status description shows any other state other than "running" or the system status description shows impaired, Auto Scaling considers the instance to be unhealthy. Thus, it terminates the instance and launches a replacement.

NEW QUESTION 252
- (Topic 3)
A user wants to find the particular error that occurred on a certain date in the AWS MySQL RDS DB. Which of the below mentioned activities may help the user to get the data easily?

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

Answer: C

Explanation:

The user can view, download, and watch the database logs using the Amazon RDS console, the Command Line Interface (CLI. or the Amazon RDS API. For the

MySQL RDS, the user can view the error log, slow query log, and general logs. The user can also view the MySQL logs easily by directing the logs to a database table in the main database and querying that table.

NEW QUESTION 254
- (Topic 3)
A user is collecting 1000 records per second. The user wants to send the data to CloudWatch using the custom namespace. Which of the below mentioned options is recommended for this activity?

A. Aggregate the data with statistics, such as Min, max, Average, Sum and Sample data and send the data to CloudWatch
B. Send all the data values to CloudWatch in a single command by separating them with a comm
C. CloudWatch will parse automatically
D. Create one csv file of all the data and send a single file to CloudWatch
E. It is not possible to send all the data in one cal
F. Thus, it should be sent one by on
G. CloudWatch will aggregate the data automatically

Answer: A

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 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. It is recommended that when the user is having multiple data points per minute, he should aggregate the data so that it will minimize the number of calls to put-metric-data. In this case it will be single call to CloudWatch instead of 1000 calls if the data is aggregated.

NEW QUESTION 255
- (Topic 3)
Which services allow the customer to retain run administrative privileges or the undertying EC2 instances? Choose 2 answers

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

Answer: AB

NEW QUESTION 257
- (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 258
- (Topic 3)
A user has configured an EC2 instance in the US-East-1a zone. The user has enabled detailed monitoring of the instance. The user is trying to get the data from CloudWatch using a CLI. Which of the below mentioned CloudWatch endpoint URLs should the user use?

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

Answer: A

Explanation:

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

NEW QUESTION 260
- (Topic 3)
A root account owner has given full access of his S3 bucket to one of the IAM users using the bucket ACL. When the IAM user logs in to the S3 console, which actions can he perform?

A. He can just view the content of the bucket

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

Answer: C

Explanation:

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

NEW QUESTION 264
- (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 268
- (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 273
- (Topic 3)
A user is planning to set up the Multi AZ feature of RDS. Which of the below mentioned conditions won't take advantage of the Multi AZ feature?

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

Answer: C

Explanation:

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

NEW QUESTION 276
- (Topic 3)
The compliance department within your multi-national organization requires that all data for your customers that reside in the European Union (EU) must not leave the EU and also
data for customers that reside in the US must not leave the US without explicit authorization.
What must you do to comply with this requirement for a web based profile management application running on EC2?

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

Answer: C

NEW QUESTION 281
- (Topic 3)
A user has launched an EBS backed EC2 instance in the US-East-1a region. The user stopped the instance and started it back after 20 days. AWS throws up an 'InsufficientInstanceCapacity' error. What can be the possible reason for this?

A. AWS does not have sufficient capacity in that availability zone
B. AWS zone mapping is changed for that user account
C. There is some issue with the host capacity on which the instance is launched
D. The user account has reached the maximum EC2 instance limit

Answer: A

Explanation:

When the user gets an 'InsufficientInstanceCapacity' error while launching or starting an EC2 instance, it means that AWS does not currently have enough available capacity to service the user request. If the user is requesting a large number of instances, there might not be enough server capacity to host them. The user can either try again later, by specifying a smaller number of instances or changing the availability zone if launching a fresh instance.

NEW QUESTION 285
- (Topic 3)
You run a web application with the following components Elastic Load Balancer (EL8), 3 Web/Application servers, 1 MySQL RDS database with read replicas, and Amazon Simple Storage Service (Amazon S3) for static content. Average response time for users is increasing slowly.
What three CloudWatch RDS metrics will allow you to identify if the database is the bottleneck? Choose 3 answers

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

Answer: ABD

NEW QUESTION 290
- (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 293
- (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 297
- (Topic 3)
An organization has configured Auto Scaling for hosting their application. The system admin wants to
understand the Auto Scaling health check process. If the instance is unhealthy, Auto Scaling launches an
instance and terminates the unhealthy instance. What is the order execution?

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

Answer: D

Explanation:

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

NEW QUESTION 298
- (Topic 3)
An organization is planning to create a user with IAM. They are trying to understand the limitations of IAM so that they can plan accordingly. Which of the below
mentioned statements is not true with respect to the
limitations of IAM?

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

Answer: A

Explanation:

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

NEW QUESTION 301
- (Topic 3)
Your business is building a new application that will store its entire customer database on a RDS MySQL database, and will have various applications and users
that will query that data for different purposes.
Large analytics jobs on the database are likely to cause other applications to not be able to get the query results they need to, before time out. Also, as your data
grows, these analytics jobs will start to take more time, increasing the negative effect on the other applications.
How do you solve the contention issues between these different workloads on the same data?

A. Enable Multi-AZ mode on the RDS instance
B. Use ElastiCache to offload the analytics job data
C. Create RDS Read-Replicas for the analytics work
D. Run the RDS instance on the largest size possible

Answer: B

NEW QUESTION 305
- (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 307
- (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 312
- (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 314
- (Topic 3)
A user has created a VPC with the public subnet. The user has created a security group for that VPC. Which of the below mentioned statements is true when a security group is created?

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

Answer: C

Explanation:

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

NEW QUESTION 316
- (Topic 3)
You have private video content in S3 that you want to serve to subscribed users on the Internet. User IDs, credentials, and subscriptions are stored in an Amazon RDS database.
Which configuration will allow you to securely serve private content to your users?

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

Answer: C

Explanation:
Reference:

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

NEW QUESTION 318
- (Topic 3)
A root account owner is trying to understand the S3 bucket ACL. Which of the below mentioned options cannot be used to grant ACL on the object using the authorized predefined group?

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

Answer: D

Explanation:

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

NEW QUESTION 320
- (Topic 3)
A user has launched an RDS postgreSQL DB with AWS. The user did not specify the maintenance window during creation. The user has configured RDS to update the DB instance type from micro to large. If the user wants to have it during the maintenance window, what will AWS do?

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

Answer: B

Explanation:

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

NEW QUESTION 324
- (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 328
- (Topic 3)
A user has setup a custom application which generates a number in decimals. The user wants to track that number and setup the alarm whenever the number is above a certain limit. The application is sending the data to CloudWatch at regular intervals for this purpose. Which of the below mentioned statements is not true with respect to the above scenario?

A. The user can get the aggregate data of the numbers generated over a minute and send it to CloudWatch
B. The user has to supply the timezone with each data point
C. CloudWatch will not truncate the number until it has an exponent larger than 126 (i.
D. (1 x 10^126.
E. The user can create a file in the JSON format with the metric name and value and supply it to CloudWatch

Answer: B

NEW QUESTION 331
- (Topic 3)
Amazon EBS snapshots have which of the following two characteristics? (Choose 2.) Choose 2 answers

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

Answer: AD

NEW QUESTION 333
- (Topic 3)
A user has launched an EC2 instance from an instance store backed AMI. The infrastructure team wants to create an AMI from the running instance. Which of the below mentioned credentials is not required while creating the AMI?

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

Answer: C

Explanation:

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

NEW QUESTION 337
- (Topic 3)
A user has created a mobile application which makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK and root account access/secret access key to connect to DynamoDB from mobile. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

A. The user should create a separate IAM user for each mobile application and provide DynamoDB access with it
B. The user should create an IAM role with DynamoDB and EC2 acces
C. Attach the role with EC2 and route all calls from the mobile through EC2
D. The application should use an IAM role with web identity federation which validates calls to DynamoDB with identity providers, such as Google, Amazon, and Facebook
E. Create an IAM Role with DynamoDB access and attach it with the mobile application

Answer: C

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. If the user is creating an app that runs on a mobile phone and makes requests to AWS, the user should not create an IAMuser and distribute the user's access key with the app. Instead, he should use an identity provider, such as Login with Amazon, Facebook, or Google to authenticate the users, and then use that identity to get temporary security credentials.

NEW QUESTION 339
- (Topic 3)
A user has created a Cloudformation stack. The stack creates AWS services, such as EC2 instances, ELB, AutoScaling, and RDS. While creating the stack it created EC2, ELB and AutoScaling but failed to create RDS. What will Cloudformation do in this scenario?

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

Answer: C

Explanation:

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

NEW QUESTION 340
- (Topic 3)
A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console?

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

Answer: D

Explanation:

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

NEW QUESTION 344
- (Topic 3)
A sys admin is planning to subscribe to the RDS event notifications. For which of the below mentioned source categories the subscription cannot be configured?

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

Answer: C

Explanation:

Amazon RDS uses the Amazon Simple Notification Service (SNS. to provide a notification when an Amazon RDS event occurs. These events can be configured for source categories, such as DB instance, DB security group, DB snapshot and DB parameter group.

NEW QUESTION 345
- (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 349
- (Topic 3)
A user is receiving a notification from the RDS DB whenever there is a change in the DB security group. The user does not want to receive these notifications for only a month. Thus, he does not want to delete the notification. How can the user configure this?

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

Answer: D

Explanation:

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

NEW QUESTION 353
- (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 357

- (Topic 3)
A user has launched an EC2 Windows instance from an instance store backed AMI. The user has also set the Instance initiated shutdown behavior to stop. What will happen when the user shuts down the OS?

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

Answer: B

Explanation:

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

NEW QUESTION 362
- (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 366
- (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.
Which of the below mentioned entries are required in the main route table to allow the instances in VPC to communicate with each other?

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

Answer: D

NEW QUESTION 371
- (Topic 3)
A user has setup an EBS backed instance and attached 2 EBS volumes to it. The user has setup a CloudWatch alarm on each volume for the disk data. The user has stopped the EC2 instance and detached the EBS volumes. What will be the status of the alarms on the EBS volume?

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

Answer: B

Explanation:

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

NEW QUESTION 373
- (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 374
- (Topic 3)
A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling AddToLoadBalancer
(which adds instances to the load balancer. process for a while. What will happen to the instances launched during the suspension period?

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

Answer: A

Explanation:

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

NEW QUESTION 378
- (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 379
- (Topic 3)
A user is running a batch process on EBS backed EC2 instances. The batch process starts a few instances to process hadoop Map reduce jobs which can run between 50 – 600 minutes or sometimes for more time. The user wants to configure that the instance gets terminated only when the process is completed. How can the user configure this with CloudWatch?

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

Answer: D

Explanation:

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

NEW QUESTION 380
- (Topic 3)
A user has created a VPC with two subnets: one public and one private. The user is planning to run the patch update for the instances in the private subnet. How can the instances in the private subnet connect to the internet?

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

Answer: D

Explanation:

A Virtual Private Cloud (VPC. is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. If the user has created two subnets (one private and one public., he would need a Network Address Translation (NAT. instance with the elastic IP address. This enables the instances in the private subnet to send requests to the internet (for example, to perform software updates..

NEW QUESTION 384
- (Topic 3)
A user is trying to connect to a running EC2 instance using SSH. However, the user gets a Host key not found error. Which of the below mentioned options is a possible reason for rejection?

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

Answer: A

Explanation:

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

NEW QUESTION 386
......

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

    All transactions are protected by VeriSign!

100% Pass Your AWS-SysOps Exam with Our Prep Materials Via below:

https://www.certleader.com/AWS-SysOps-dumps.html