SOA-C01 Dumps

AWS Certified SysOps Administrator - Associate

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

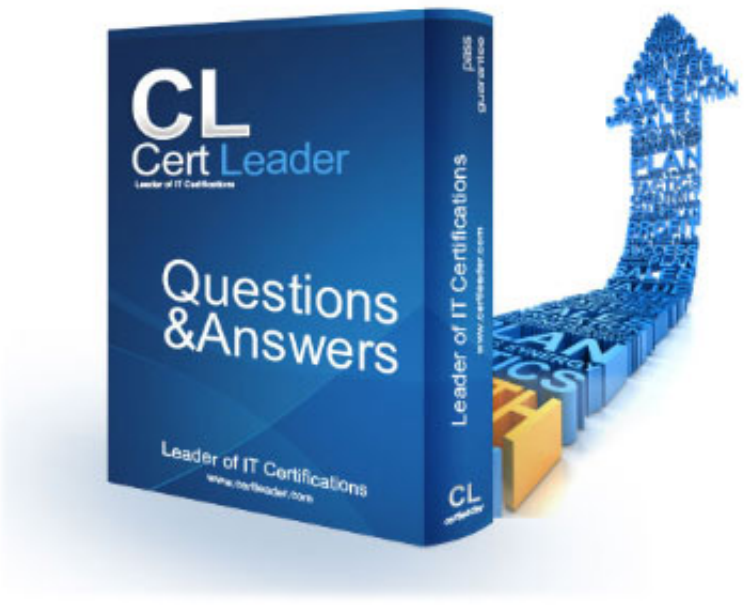

NEW QUESTION 1
When preparing for a compliance assessment of your system built inside of AWS. what are three best-practices for you to prepare for an audit?
Choose 3 answers

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

Answer: ABD

NEW QUESTION 2
You have 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: D

Explanation:
The question is about most ??scalable solution for communicating?? for SQS that is parallel processing of SQS messages.
See also:
?V https://aws.amazon.com/articles/1464
?Vhttp://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/throughput.html

NEW QUESTION 3
An application that you are managing has EC2 instances & Dynamo OB tables deployed to several AWS Regions. In order to monitor the performance of the application globally, you would like to see two graphs: 1) Avg CPU Utilization across all EC2 instances and 2) Number of Throttled Requests for all DynamoDB tables.
How can you accomplish this?

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

Answer: A

Explanation:
Correct answer should be A. When you turn on detailed monitoring in CloudWatch, you can get 1) Avg CPU Utilization across all EC2 instances and 2) Number of Throttled Requests for all DynamoDB tables
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/GetSingleMetricAllDimen sions.html

NEW QUESTION 4
Your entire AWS infrastructure lives inside of one Amazon VPC You have an Infrastructure monitoring application running on an Amazon instance in Availability Zone (AZ) A of the region, and another application instance running in AZ B. The monitoring application needs to make use of ICMP ping to confirm network reachability of the instance hosting the application.
Can you configure the security groups for these instances to only allow the ICMP ping to pass from the monitoring instance to the application instance and nothing else" If so how?

A. N
B. Two instances in two different AZ's can't talk directly to each other via ICMP ping as that protocol is not allowed across subnet (i.e., broadcast) boundaries
C. Ye
D. Both the monitoring instance and the application instance have to be a part of the same security group, and that security group needs to allow inbound ICMP
E. Ye
F. The security group for the monitoring instance needs to allow outbound ICMP and the application instance's security group needs to allow Inbound ICMP
G. Yes, Both the monitoring instance's security group and the application instance's security group need to allow both inbound and outbound ICMP ping packets since ICMP is not a connection- oriented protocol

Answer: C

NEW QUESTION 5
Which services allow the customer to retain full administrative privileges of the underlying EC2 instances?
Choose 2 answers

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

Answer: AC

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

NEW QUESTION 6
You are designing a system that has a Bastion host. This component needs to be highly available without human intervention.
Which of the following approaches would you select?

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

Answer: C

NEW QUESTION 7
Which of the following requires a custom CloudWatch metric to monitor?

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

Answer: C

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

NEW QUESTION 8
When creation of an EBS snapshot Is initiated but not completed the EBS volume?

A. Cannot be 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:
Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html

NEW QUESTION 9
You are using ElastiCache Memcached to store session state and cache database queries in your infrastructure. You notice in CloudWatch that Evictions and GetMisses are Doth very high.
What two actions could you take to rectify this? Choose 2 answers

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

Answer: AB

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

NEW QUESTION 10
You are running a database on an EC2 instance, with the data stored on Elastic Block Store (EBS) for persistence. At times throughout the day, you are seeing large variance in the response times of the database queries Looking into the instance with the isolate command you see a lot of wait time on the disk volume that the database's data is stored on.
What two ways can you improve the performance of the database's storage while maintaining the current persistence of the data?
Choose 2 answers

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

Answer: BC

NEW QUESTION 10
You are attempting to connect to an instance in Amazon VPC without success. You have already verified that the VPC has an Internet Gateway (IGW) the instance has an associated Elastic IP (EIP) and correct security group rules are in place.
Which VPC component should you evaluate next?

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

Answer: B

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

NEW QUESTION 12
You are tasked with the migration of a highly trafficked Node JS application to AWS In order to comply with organizational standards Chef recipes must be used to configure the application servers that host this application and to support application lifecycle events.
Which deployment option meets these requirements while minimizing administrative burden?

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

Answer: A

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

NEW QUESTION 16
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 18
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:
https://aws.amazon.com/blogs/aws/amazon-s3-multipart-upload/

NEW QUESTION 20
You are running a web-application on AWS consisting of the following components an Elastic Load Balancer (ELB) an Auto-Scaling Group of EC2 instances running Linux/PHP/Apache, and Relational Database Service (RDS) MySQL.
Which security measures fall into AWS's responsibility?

A. Protect the EC2 instances against unsolicited access by enforcing the principle of least-privilege access
B. Protect against IP spoofing or packet sniffing
C. Assure all communication between EC2 instances and ELB is encrypted
D. Install latest security patches on EL
E. RDS and EC2 instances

Answer: B

NEW QUESTION 25
When an EC2 EBS-backed (EBS root) instance is stopped, what happens to the data on any ephemeral store volumes?

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

Answer: A

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

NEW QUESTION 30
If you want to launch Amazon Elastic Compute Cloud (EC2) Instances and assign each Instance a predetermined private IP address you should:

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

Answer: C

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

NEW QUESTION 32
What would happen to an RDS (Relational Database Service) multi-Availability Zone deployment of the primary OB instance fails?

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

Answer: D

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

NEW QUESTION 37
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 39
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 41
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 46
A user has developed an application which is required to send the data to a NoSQL database. The user wants to decouple the data sending such that the application keeps processing and sending data but does not wait for an acknowledgement of DB. Which of the below mentioned applications helps in this scenario?

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

Answer: C

Explanation:
Amazon Simple Queue Service (SQS. is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. In this case, the user can use AWS SQS to send messages which are received from an application and sent to DB. The application can continue processing data without waiting for any acknowledgement from DB. The user can use SQS to transmit any volume of data without losing messages or requiring other services to always be available.

NEW QUESTION 51
An organization has created 50 IAM users. The organization has introduced a new policy which will change the access of an IAM user. How can the organization implement this effectively so that there is no need to apply the policy at the individual user level?

A. Use the IAM groups and add users as per their role to different groups and apply policy to group
B. The user can create a policy and apply it to multiple users in a single go with the AWS CLI
C. Add each user to the IAM role as per their organization role to achieve effective policy setup
D. Use the IAM role and implement access at the role level

Answer: A

Explanation:
With AWS IAM, a group is a collection of IAM users. A group allows the user to specify permissions for a collection of users, which can make it easier to manage the permissions for those users. A group helps an organization manage access in a better way; instead of applying at the individual level, the organization can apply at the group level which is applicable to all the users who are a part of that group.

NEW QUESTION 55
A user is planning to use AWS Cloud formation for his automatic deployment requirements. Which of the below mentioned components are required as a part of the template?

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

Answer: D

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

NEW QUESTION 58
A user has recently started using EC2. The user launched one EC2 instance in the default subnet in EC2-VPC Which of the below mentioned options is not attached or available with the EC2 instance when it is launched?

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

Answer: C

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

NEW QUESTION 61
A user has launched an EC2 instance. The user is planning to setup the CloudWatch alarm. Which of the below mentioned actions is not supported by the CloudWatch alarm?

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

Answer: A

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

NEW QUESTION 65
An organization is planning to create 5 different AWS accounts considering various security requirements. The organization wants to use a single payee account by using the consolidated billing option. Which of the below mentioned statements is true with respect to the above information?

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

Answer: B

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

NEW QUESTION 70
A user has created a web application with Auto Scaling. The user is regularly monitoring the application and he observed that the traffic is highest on Thursday and Friday between 8 AM to 6 PM. What is the best solution to handle scaling in this case?

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

Answer: B

Explanation:
Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. In this case the load increases by Thursday and decreases by Friday. Thus, the user can setup the scaling activity based on the predictable traffic patterns of the web application using Auto Scaling scale by Schedule.
http://docs.aws.amazon.com/cli/latest/reference/opsworks/set-time-based-auto-scaling.html

NEW QUESTION 72
A user has setup a CloudWatch alarm on an EC2 action when the CPU utilization is above 75%. The alarm sends a notification to SNS on the alarm state. If the user wants to simulate the alarm action how can he achieve this?

A. Run activities on the CPU such that its utilization reaches above 75%
B. From the AWS console change the state to ??Alarm??
C. The user can set the alarm state to ??Alarm?? using CLI
D. Run the SNS action manually

Answer: C

Explanation:
Amazon CloudWatch alarms watch a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods.The user can test an alarm by setting it to any state using the SetAlarmState API (mon-set-alarm-state command.. This temporary state change lasts only until the next alarm comparison occurs.

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

NEW QUESTION 74
A user is trying to setup a scheduled scaling activity using Auto Scaling. The user wants to setup the recurring schedule. Which of the below mentioned parameters is not required in this case?

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

Answer: A

Explanation:
Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. The user can also configure the recurring schedule action which will follow the Linux cron format. If the user is setting a recurring event, it is required that the user specifies the Recurrence value (in a cron format., end time (not compulsory but recurrence will stop after this. and the Auto Scaling group for which the scaling activity is to be scheduled.

NEW QUESTION 78
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 requirement 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 82
A user has created a subnet with VPC and launched an EC2 instance in that subnet with only default settings. Which of the below mentioned options is ready to use on the EC2 instance as soon as it is launched?

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

Answer: B

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

NEW QUESTION 87
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 90
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 95
An organization has created 50 IAM users. The organization wants that each user can change their password but cannot change their access keys. How can the organization achieve this?

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

Answer: D

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

NEW QUESTION 98
A user has created a photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly. Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

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

Answer: D

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

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

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

Answer: D

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

NEW QUESTION 107
A user is accessing RDS from an application. The user has enabled the Multi AZ feature with the MS SQL RDS DB. During a planned outage how will AWS ensure that a switch from DB to a standby replica will not affect access to the application?

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

Answer: B

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

NEW QUESTION 108
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 110
A user has launched an EBS backed instance. The user started the instance at 9 AM in the morning. Between 9 AM to 10 AM, the user is testing some script. Thus, he stopped the instance twice and restarted it. In the same hour the user rebooted the instance once. For how many instance hours will AWS charge the user?

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

Answer: A

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

NEW QUESTION 115
A user has launched a large EBS backed EC2 instance in the US-East-1a region. The user wants to achieve Disaster Recovery (DR. for that instance by creating another small instance in Europe. How can the user achieve DR?

A. Copy the running instance using the ??Instance Copy?? command to the EU region
B. Create an AMI of the instance and copy the AMI to the EU regio
C. Then launch the instance from the EU AMI
D. Copy the instance from the US East region to the EU region
E. Use the ??Launch more like this?? option to copy the instance from one region to another

Answer: B

Explanation:
To launch an EC2 instance it is required to have an AMI in that region. If the AMI is not available in that region, then create a new AMI or use the copy command to copy the AMI from one region to the other region.

NEW QUESTION 120
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 125
An organization is using cost allocation tags to find the cost distribution of different departments and projects. One of the instances has two separate tags with the key/ value as ??InstanceName/HR??, ??CostCenter/HR??. What will AWS do in this case?

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

Answer: D

Explanation:

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

NEW QUESTION 127
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 130
A user is launching an EC2 instance in the US East region. Which of the below mentioned options is recommended by AWS with respect to the selection of the availability zone?

A. Always select the US-East-1-a zone for HA
B. Do not select the AZ; instead let AWS select the AZ
C. The user can never select the availability zone while launching an instance
D. Always select the AZ while launching an instance

Answer: B

Explanation:
When launching an instance with EC2, AWS recommends not to select the availability zone (AZ. AWS specifies that the default Availability Zone should be accepted. This is because it enables AWS to select the best Availability Zone based on the system health and available capacity. If the user launches additional instances, only then an Availability Zone should be specified. This is to specify the same or different AZ from the running instances.

NEW QUESTION 131
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 136
A user is trying to connect to a running EC2 instance using SSH. However, the user gets a connection time out error. Which of the below mentioned options is not a possible reason for rejection?

A. The access key to connect to the instance is wrong
B. The security group is not configured properly
C. The private key used to launch the instance is not correct
D. The instance CPU is heavily loaded

Answer: A

Explanation:
If the user is trying to connect to a Linux EC2 instance and receives the connection time out error the probable reasons are:
Security group is not configured with the SSH port The private key pair is not right
The user name to login is wrong
The instance CPU is heavily loaded, so it does not allow more connections

NEW QUESTION 139
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 142
A sys admin is trying to understand EBS snapshots. Which of the below mentioned statements will not be useful to the admin to understand the concepts about a snapshot?

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

Answer: A

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

NEW QUESTION 146
A user has launched two EBS backed EC2 instances in the US-East-1a region. The user wants to change the zone of one of the instances. How can the user change it?

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

Answer: D

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

NEW QUESTION 149
An organization has added 3 of his AWS accounts to consolidated billing. One of the AWS accounts has purchased a Reserved Instance (RI. of a small instance size in the US-East-1a zone. All other AWS accounts are running instances of a small size in the same zone. What will happen in this case for the RI pricing?

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

Answer: C

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

NEW QUESTION 152
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 155
A user is running one instance for only 3 hours every day. The user wants to save some cost with the instance. Which of the below mentioned Reserved Instance categories is advised in this case?

A. The user should not use RI; instead only go with the on-demand pricing

B. The user should use the AWS high utilized RI
C. The user should use the AWS medium utilized RI
D. The user should use the AWS low utilized RI

Answer: A

Explanation:
The AWS Reserved Instance provides the user with an option to save some money by paying a one- time fixed amount and then save on the hourly rate. It is advisable that if the user is having 30% or more usage of an instance per day, he should go for a RI. If the user is going to use an EC2 instance for more than 2200-2500 hours per year, RI will help the user save some cost. Here, the instance is not going to run for less than 1500 hours. Thus, it is advisable that the user should use the on- demand pricing.

NEW QUESTION 158
A user has setup an RDS DB with Oracle. The user wants to get notifications when someone modifies the security group of that DB. How can the user configure that?

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

Answer: C

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

NEW QUESTION 159
A user has created an ELB with Auto Scaling. Which of the below mentioned offerings from ELB helps the user to stop sending new requests traffic from the load balancer to the EC2 instance when the instance is being deregistered while continuing in-flight requests?

A. ELB sticky session
B. ELB deregistration check
C. ELB connection draining
D. ELB auto registration Off

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.

NEW QUESTION 163
An AWS root account owner is trying to create a policy to access RDS. Which of the below mentioned statements is true with respect to the above information?

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

Answer: C

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the account owner wants to create a policy for RDS, the owner has to create an IAM user and define the policy which entitles the IAM user with various RDS services such as Launch Instance, Manage security group, Manage parameter group etc.

NEW QUESTION 166
A user has configured an SSL listener at ELB as well as on the back-end instances. Which of the below mentioned statements helps the user understand ELB traffic handling with respect to the SSL listener?

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

Answer: D

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

NEW QUESTION 167
A user has created a Cloudformation stack. The stack creates AWS services, such as EC2 instances, ELB, AutoScaling, and RDS. While creating the stack it created EC2, ELB and AutoScaling but failed to

create RDS. What will Cloudformation do in this scenario?

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

Answer: C

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

NEW QUESTION 169
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 174
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 175
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 179
A sysadmin has created the below mentioned policy on an S3 bucket named cloudacademy. What does this policy define?
"Statement": [{
"Sid": "Stmt1388811069831",
"Effect": "Allow", "Principal": { "AWS": "*"},
"Action": ["s3:GetObjectAcl", "s3:ListBucket"], "Resource": ["arn:aws:s3:::cloudacademy]
}]

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

Answer: D

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

NEW QUESTION 183
A user has launched two EBS backed EC2 instances in the US-East-1a region. The user wants to change the zone of one of the instances. How can the user change it?

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

Answer: B

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

NEW QUESTION 188
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 191
A sys admin has enabled logging on ELB. Which of the below mentioned fields will not be a part of the log file name?

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

Answer: B

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

NEW QUESTION 194
A user has created a queue named ??awsmodule?? with SQS. One of the consumers of queue is down for 3 days and then becomes available. Will that component receive message from queue?

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

Answer: A

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

NEW QUESTION 199

An organization has created one IAM user and applied the below mentioned policy to the user. What entitlements do the IAM users avail with this policy?
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "ec2:Describe*", "Resource": "*"
},
{
"Effect": "Allow"
"Action": ["cloudwatch:ListMetrics", "cloudwatch:GetMetricStatistics", "cloudwatch:Describe*"
],
"Resource": "*"
},
{
"Effect": "Allow",
"Action": "autoscaling:Describe*", "Resource": "*"
}
]
}

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

Answer: D

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If an organization wants to setup read only access to EC2 for a particular user, they should mention the action in the IAM policy which entitles the user for Describe rights for EC2, CloudWatch, Auto Scaling and ELB. In the policy shown below, the user will have read only access for EC2 and EBS, CloudWatch and Auto Scaling. Since ELB is not mentioned as a
part of the list, the user will not have access to ELB.
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "ec2:Describe*", "Resource": "*"
},
{
"Effect": "Allow", "Action": ["cloudwatch:ListMetrics",
"cloudwatch:GetMetricStatistics", "cloudwatch:Describe*"
],
"Resource": "*"
},
{
"Effect": "Allow",
"Action": "autoscaling:Describe*", "Resource": "*"
}
]
}

NEW QUESTION 203
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:
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
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 209
George has shared an EC2 AMI created in the US East region from his AWS account with Stefano. George copies the same AMI to the US West region. Can Stefano access the copied AMI of George??s account from the US West region?

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

Answer: A

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

NEW QUESTION 214
A user has created a VPC with a subnet and a security group. The user has launched an instance in that subnet and attached a public IP. The user is still unable to connect to the instance. The internet gateway has also been created. What can be the reason for the error?

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

Answer: A

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

NEW QUESTION 219
George has launched three EC2 instances inside the US-East-1a zone with his AWS account. Ray has launched two EC2 instances in the US-East-1a zone with his AWS account. Which of the below entioned statements will help George and Ray understand the availability zone (AZ. concept better?

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

Answer: D

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

NEW QUESTION 222
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 227
A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created public and VPN only subnets along with hardware VPN access to connect to the user??s data centre. The user has not yet launched any instance as well as modified or deleted any setup. He wants to delete this VPC from the console. Will the console allow the user to delete the VPC?

A. Yes, the console will delete all the setups and also delete the virtual private gateway
B. No, the console will ask the user to manually detach the virtual private gateway first and then allow deleting the VPC
C. Yes, the console will delete all the setups and detach the virtual private gateway
D. No, since the NAT instance is running

Answer: C

Explanation:
The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all traffic of the VPN subnet. If the virtual private gateway is attached with VPC and the user deletes the VPC from the console it will first detach the gateway automatically and only then delete the VPC.

NEW QUESTION 232
A user is trying to create a PIOPS EBS volume with 4000 IOPS and 100 GB size. AWS does not allow the user to create this volume. What is the possible root cause for this?

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

Answer: A

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

NEW QUESTION 235
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 236
A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24. What will happen in this scenario?

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

Answer: D

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

NEW QUESTION 237

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

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

Answer: B

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

NEW QUESTION 242
A sys admin is using server side encryption with AWS S3. Which of the below mentioned statements helps the user understand the S3 encryption functionality?

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

Answer: A

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

NEW QUESTION 246
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 251
A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created a public subnet CIDR (20.0.0.0/24. and VPN only subnets CIDR (20.0.1.0/24. along with the VPN gateway (vgw-12345. to connect to the user??s data centre. The user??s data centre has CIDR 172.28.0.0/12. The user has also setup a NAT instance (i-123456. to allow traffic to the internet from the VPN subnet. Which of the below mentioned options is not a valid entry for the main route table in this scenario?

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

Answer: A

Explanation:
The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all traffic of the VPN subnet. If the user has setup a NAT instance to route all the internet requests then all requests to the internet should be routed to it. All requests to the organization??s DC will be routed to the VPN gateway.
Here are the valid entries for the main route table in this scenario:
Destination: 0.0.0.0/0 & Target: i-12345 (To route all internet traffic to the NAT Instance.
Destination: 172.28.0.0/12 & Target: vgw-12345 (To route all the organization??s data centre traffic to the VPN gateway.
Destination: 20.0.0.0/16 & Target: local (To allow local routing in VPC.

NEW QUESTION 254
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?Vspecific XML schema. The user cannot grant permissions to other users (IAM users. in his account.

NEW QUESTION 258
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 260
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 261
A user is using the AWS SQS to decouple the services. Which of the below mentioned operations is not supported by SQS?

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

Answer: D

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

NEW QUESTION 264
A user has launched an EC2 instance. However, due to some reason the instance was terminated. If the user wants to find out the reason for termination, where can he find the details?

A. It is not possible to find the details after the instance is terminated
B. The user can get information from the AWS console, by checking the Instance description under the State transition reason label
C. The user can get information from the AWS console, by checking the Instance description under the Instance Status Change reason label
D. The user can get information from the AWS console, by checking the Instance description under the Instance Termination reason label

Answer: D

Explanation:
An EC2 instance, once terminated, may be available in the AWS console for a while after termination. The user can find the details about the termination from the description tab under the label State transition reason. If the instance is still running, there will be no reason listed. If the user has explicitly stopped or terminated the instance, the reason will be ??User initiated shutdown??.

NEW QUESTION 267
An organization (Account ID 123412341234. has attached the below mentioned IAM policy to a user. What does this policy statement entitle the user to perform?
{
"Version": "2012-10-17",
"Statement": [{

"Sid": "AllowUsersAllActionsForCredentials", "Effect": "Allow",
"Action": ["iam:*LoginProfile", "iam:*AccessKey*",
"iam:*SigningCertificate*"
],
"Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
}]
}

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

Answer: D

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

NEW QUESTION 270
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 274
A user has provisioned 2000 IOPS to the EBS volume. The application hosted on that EBS is experiencing less IOPS than provisioned. Which of the below mentioned options does not affect the IOPS of the volume?

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

Answer: D

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

NEW QUESTION 279
A storage admin wants to encrypt all the objects stored in S3 using server side encryption. The user does not want to use the AES 256 encryption key provided by S3. How can the user achieve this?

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

Answer: D

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

NEW QUESTION 284
A user is having data generated randomly based on a certain event. The user wants to upload that data to CloudWatch. It may happen that event may not have data generated for some period due to andomness. Which of the below mentioned options is a recommended option for this case?

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

Answer: C

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

NEW QUESTION 285
A user is sending the data to CloudWatch using the CloudWatch API. The user is sending data 90 minutes in the future. What will CloudWatch do in this case?

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

Answer: A

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

NEW QUESTION 288
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 290
A user has created a VPC with two subnets: one public and one private. The user is planning to run the patch update for the instances in the private subnet. How can the instances in the private subnet connect to theinternet?

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

Answer: D

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

NEW QUESTION 291
A user is running a batch process on EBS backed EC2 instances. The batch process starts a few instances to process hadoop Map reduce jobs which can run between 50 ?V 600 minutes or sometimes for more time. The user wants to configure that the instance gets terminated only when the process is completed. How can the user configure this with CloudWatch?

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

Answer: D

Explanation:

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

NEW QUESTION 296
A user has enabled versioning on an S3 bucket. The user is using server side encryption for data at rest. If the user is supplying his own keys for encryption (SSE-C., what is recommended to the user for the purpose of security?

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

Answer: D

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

NEW QUESTION 301
A user has launched an EC2 Windows instance from an instance store backed AMI. The user wants to convert the AMI to an EBS backed AMI. How can the user convert it?

A. Attach an EBS volume to the instance and unbundle all the AMI bundled data inside the EBS
B. A Windows based instance store backed AMI cannot be converted to an EBS backed AMI
C. It is not possible to convert an instance store backed AMI to an EBS backed AMI
D. Attach an EBS volume and use the copy command to copy all the ephermal content to the EBS Volume

Answer: B

Explanation:
Generally when a user has launched an EC2 instance from an instance store backed AMI, it can be converted to an EBS backed AMI provided the user has attached the EBS volume to the instance and
unbundles the AMI data to it. However, if the instance is a Windows instance, AWS does not allow this. In this case, since the instance is a Windows instance, the user cannot convert it to an EBS backed AMI.

NEW QUESTION 303
A sysadmin has created the below mentioned policy on an S3 bucket named cloudacademy. The bucket has both AWS.jpg and index.html objects. What does this policy define?
"Statement": [{
"Sid": "Stmt1388811069831",
"Effect": "Allow", "Principal": { "AWS": "*"},
"Action": ["s3:GetObjectAcl", "s3:ListBucket", "s3:GetObject"], "Resource": ["arn:aws:s3:::cloudacademy/*.jpg]
}]

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

Answer: B

NEW QUESTION 306
A user is trying to pre-warm a blank EBS volume attached to a Linux instance. Which of the below mentioned steps should be performed by the user?

A. There is no need to pre-warm an EBS volume
B. Contact AWS support to pre-warm
C. Unmount the volume before pre-warming
D. Format the device

Answer: C

Explanation:
When the user creates a new EBS volume or restores a volume from the snapshot, the back-end storage blocks are immediately allocated to the user EBS. However, the first time when the user is trying to access a block of the storage, it is recommended to either be wiped from the new volumes or instantiated from the snapshot (for restored volumes. before the user can access the block. This preliminary action takes time and can cause a 5 to 50 percent loss of IOPS for the volume when the block is accessed for the first time. To avoid this it is required to pre warm the volume. Pre-warming an EBS volume on a Linux instance requires that the user should unmount the blank device first and then write all the blocks on the device using a command, such as ??dd??.

NEW QUESTION 308
A user has enabled termination protection on an EC2 instance. The user has also set Instance initiated shutdown behaviour to terminate. When the user shuts down the instance from the OS, what will happen?

A. The OS will shutdown but the instance will not be terminated due to protection
B. It will terminate the instance
C. It will not allow the user to shutdown the instance from the OS

D. It is not possible to set the termination protection when an Instance initiated shutdown is set to Terminate

Answer: B

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

NEW QUESTION 310
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 311
When an EC2 instance mat is backed by an S3-Dased AMI is terminated, what happens to the data on the root volume?

A. Data is automatically deleted
B. Data is automatically saved as an EBS snapshot.
C. Data is unavailable until the instance is restarted
D. Data is automatically saved as an EBS volume.

Answer: A

NEW QUESTION 314
How can you secure data at rest on an EBS volume?

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

Answer: E

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

NEW QUESTION 315
You run a web application with the following components Elastic Load Balancer (EL8), 3 Web/Application servers, 1 MySQL RDS database with read replicas, and Amazon Simple Storage Service (Amazon S3) for static content. Average response time for users is increasing slowly.
What three CloudWatch RDS metrics will allow you to identify if the database is the bottleneck? Choose 3 answers

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

Answer: ABE

NEW QUESTION 320
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 checksum.
C. A success code is inserted into the S3 object metadata.
D. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful.
E. Amazon S3 is engineered for 99.999999999% durabilit
F. Therefore, there is no need to confirm that data was inserted.

Answer: C

Explanation:
To ensure that data is not corrupted traversing the network, use the Content-MD5 form field. When you use this form field, Amazon S3 checks the object against the provided MD5 value. If they do not match, Amazon S3 returns an error.
success_action_status
The status code returned to the client upon successful upload if success_action_redirectis not specified.
Accepts the values 200, 201, or 204 (default).

If the value is set to 200 or 204, Amazon S3 returns an empty document with a 200 or 204 status code.
If the value is set to 201, Amazon S3 returns an XML document with a 201 status code.
If the value is not set or if it is set to an invalid value, Amazon S3 returns an empty document with a 204 status code.
Type: String Default: None Note
Some versions of the Adobe Flash player do not properly handle HTTP responses with an empty body. To support uploads through Adobe Flash, we recommend setting success_action_statusto 201.
Source: http://docs.aws.amazon.com/AmazonS3/latest/API/RESTObjectPOST.html

NEW QUESTION 321
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 324
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: D

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

NEW QUESTION 328
A .NET application that you manage is running in Elastic Beanstalk. Your developers tell you they will need access to application log files to debug issues that arise. The infrastructure will scale up and down.
How can you ensure the developers will be able to access only the log files?

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

Answer: A

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

NEW QUESTION 332
Your mission is to create a lights-out datacenter environment, and you plan to use AWS OpsWorks to accomplish this. First you created a stack and added an App Server layer with an instance running in it. Next you added an application to the instance, and now you need to deploy a MySQL RDS database instance.
Which of the following answers accurately describe how to add a backend database server to an OpsWorks stack? Choose 3 answers

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

Answer: BCE

NEW QUESTION 335
A company has an AWS account that contains three VPCs (Dev, Test, and Prod) in the same region.
Test is peered to both Prod and Dev. All VPCs have non-overlapping CIDR blocks. The company wants to push minor code releases from Dev to Prod to speed up

time to market. Which of the following options helps the company accomplish this?

A. Create a new peering connection Between Prod and Dev along with appropriate routes.
B. Create a new entry to Prod in the Dev route table using the peering connection as the target.
C. Attach a second gateway to De
D. Add a new entry in the Prod route table identifying the gateway as the target.
E. The VPCs have non-overlapping CIDR blocks in the same accoun
F. The route tables contain local routes for all VPCs.

Answer: A

Explanation:
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/vpc-pg.pdf

NEW QUESTION 339
An instance is launched into a VPC subnet with the network ACL configured to allow all inbound traffic and deny all outbound traffic. The instance's security group is configured to allow SSH from any IP address and deny all outbound traffic. What changes need to be made to allow SSH access to the instance?

A. The outbound security group needs to be modified to allow outbound traffic.
B. The outbound network ACL needs to be modified to allow outbound traffic.
C. Nothing, it can be accessed from any IP address using SSH.
D. Both the outbound security group and outbound network ACL need to be modified to allow outbound traffic.

Answer: B

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

NEW QUESTION 343
Which of the following are true regarding encrypted Amazon Elastic Block Store (EBS) volumes? Choose 2 answers

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

Answer: AB

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

NEW QUESTION 346
A photo-sharing service stores pictures in Amazon Simple Storage Service (S3) and allows application sign-in using an OpenID Connect-compatible identity provider. Which AWS Security Token Service approach to temporary access should you use for the Amazon S3 operations?

A. SAML-based Identity Federation
B. Cross-Account Access
C. AWS Identity and Access Management roles
D. Web Identity Federation

Answer: D

NEW QUESTION 350
An Auto-Scaling group spans 3 AZs and currently has 4 running EC2 instances. When Auto Scaling needs to terminate an EC2 instance by default, AutoScaling will:
Choose 2 answers

A. Allow at least five minutes for Windows/Linux shutdown scripts to complete, before terminating the instance.
B. Terminate the instance with the least active network connection
C. If multiple instances meet this criterion, one will be randomly selected.
D. Send an SNS notification, if configured to do so.
E. Terminate an instance in the AZ which currently has 2 running EC2 instances.
F. Randomly select one of the 3 AZs, and then terminate an instance in that AZ.

Answer: CD

Explanation:
http://docs.aws.amazon.com/autoscaling/latest/userguide/as-instance-termination.html

NEW QUESTION 355
A company uses AWS Organization with a multi-account structure. A Syslog Administrator was notified that an IAM user with the System Administrator policy applied was not able to launch any Amazon EC2 instance using a public?

Why is this occurring?

A. The account is an AWS Organization master account, and by default it cannot provision EC2 instances.
B. The account is an AWS Organization member account, and a service control policy is denying provisioning of EC2 instances.
C. The account AWS Organization master account, and it does not have an access key activated for the IAM account.
D. The account is an AWS Organization master account, and it does not have an access key activated for the IAM account.

Answer: B

Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html

NEW QUESTION 357
A company Development team to access the AWS Management Console. A System Administrator has been asked to find a solution so that the Developers can sign in to the console using Active Directory (AD) credentials and not as IAM users.
What steps should the Systems Administrator take to enable functionality?

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

Answer: A

NEW QUESTION 359
A SysOps Administrator must take a team's single existing AWS CloudFormation template and split it into smaller, service specific template. All of the service in the template reference a single, shared Amazon S3 bucket.
What should the Administrator do to ensure that this S3 bucket can be referenced by all the service templates?

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

Answer: D

NEW QUESTION 362
A SysOps Administrator supports a legacy application that is hardcoded to service example.com. The application has recently been moved to AWS. The external DNS are managed by a third-party provider. The Administrator has set up an internal domain for example.com and configured this record using Amazon Route.
What solution offers the MOST efficient way to have instances in the same account resolve to the Route 53 service instead of the provider?

A. Hardcode the name server record to the internal Route 53 IP address tor each instance
B. Enable DNS resolution in the subnets as required
C. Ensure that DNS resolution is enabled on the VPC
D. Create an OS-specific hardcoded entry tor DNS resolution to the private URL

Answer: C

Explanation:
Using DNS with Your VPC
Domain Name System (DNS) is a standard by which names used on the Internet are resolved to their corresponding IP addresses. A DNS hostname is a name that uniquely and absolutely names a computer; it's composed of a host name and a domain name. DNS servers resolve DNS hostnames to their corresponding IP addresses.
Public IPv4 addresses enable communication over the Internet, while private IPv4 addresses enable communication within the network of the instance (either EC2-Classic or a VPC). For more information, see IP Addressing in Your VPC.
We provide an Amazon DNS server. To use your own DNS server, create a new set of DHCP options for your VPC. For more information, see DHCP Options Sets.
Contents
DNS Hostnames
DNS Support in Your VPC DNS Limits
Viewing DNS Hostnames for Your EC2 Instance Updating DNS Support for Your VPC
Using Private Hosted Zones

NEW QUESTION 364
An Amazon EC2 instance is unable to connect to an SMTP server in a different subnet. Other instances are successfully communication with the SMTP servers, however Flow Logs have been enabled on the SMTP server's network interface and show the following information

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

Answer: D

NEW QUESTION 365
An errant process is known to use in an entire processor and run at 100%. A SysOps Administrator wants to automate restarting the instance once the problem occurs for more than minutes.
How can this be accomplished?

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

Answer: D

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

NEW QUESTION 367
The Security team has decided that there will be no public internet access to HTTP (TCP port 80) because it is moving to HTTP for all incoming web traffic. The team a SysOps Administrator to provide a report on any security groups that are not compliant.
What should the SysOps Administrator do to provide near real-time compliance reporting?

A. Enable AWS Trusted Advisor and show the security team that the Security groups unrestricted access check will alarm
B. Schedule an AWS Lambda function to run hourly to scan and evaluate all security groups and send a report to the Security team
C. Use AWS Config to enable the restricted-common ports rule and add port 80 to the parameters
D. Use Amazon Inspector lo evaluate the security groups during scans and send the completed reports to the Security team

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/trustedadvisor/best-practices/

NEW QUESTION 368
A company's customers are reporting increased latency while accessing static web contact from Amazon S3. A SysOps Administrator a very high rate of read operations on a particular S3 bucket. What will minimize latency by reducing lead on the S3 bucket?

A. Migrate the S3 bucket to a region that is end users; geographic locations.
B. Use cross-region replication to replicate all the data to another region
C. Create an Amazon Cloud Front distribution with the bucket as the origin.
D. Use Amazon ElastiCache to cache data being server from Amazon S3

Answer: C

Explanation:
Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within
a developer-friendly environment. CloudFront is integrated with AWS ?V both physical locations that are directly connected to the AWS global infrastructure, as well as other AWS services. CloudFront works seamlessly with services including AWS Shield for DDoS mitigation, Amazon S3, Elastic Load Balancing or Amazon EC2 as origins for your applications, and Lambda@Edge to run custom code closer to customers?? users and to customize the user experience. You can get started with the Content Delivery Network in minutes, using the same AWS tools that you're already familiar with: APIs, AWS Management Console, AWS CloudFormation, CLIs, and SDKs. Amazon's CDN offers a simple, pay-as-you-go pricing model with no upfront fees or required long-term contracts, and support for the CDN is included in your existing AWS Support subscription.

NEW QUESTION 369
After a particularly high bill, an organization wants to review the use of AWS services.
What AWS service will allow the SysOps Administrator to quickly view this information to shared it, and will also forest expenses for the billing period?

A. AWS Trusted Advisor
B. Amazon QuickSight
C. AWS Cost and Usage Report
D. AWS Cost Explorer

Answer: C

NEW QUESTION 371
A company application stores document within an Amazon S3 bucket. The application is running on Amazon EC3 in a VPC. A recent change in security requirement states traffic between the company's application and the S3 bucket must leave the Amazon network.
What AWS feature can provide this functionality?

A. Security groups
B. NAT gateways
C. Virtual private gateway
D. Gateway VPC endpoint

Answer: D

Explanation:
A VPC endpoint enables you to create a private connection between your VPC and another AWS service without requiring access over the Internet, through a NAT device, a VPN connection, or AWS Direct Connect. Endpoints are virtual devices.

NEW QUESTION 373
A SysOps Administrator must monitor a fleet of Amazon EC2 Linux instance with the constraint that no agent be installed. The SysOps administrator Chooses Amazon CloudWatch as the monitoring tool.
Which metrics can be measured given the constraints? (Select THREE.)

A. CPU Utilization
B. Disk Read Operations
C. Memory Utilization
D. Network Packets in
E. Network Packets Dropped
F. CPU Ready Time

Answer: ABD

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/viewing_metrics_with_cloudwatch.html

NEW QUESTION 375
A company has a VoIP application deployed on AWS. The application is accessed by employees in a remote office and is extremely sensitive to any latency and packets loss. Minimize latency and packet loos is a higher priority than minimizing cost.
Employees are reporting occasional difficulties accessing the application. The Local Network Engineer has completed thorough troubleshooting on the LAN and unable to identify any signs of congestion or equipment failure that may be causing the issue.
What is the BEST way to address the connectivity issues between the remote office and the application?

A. Configure a VPN connection to the VPC Route all traffic to the application via the VPN connection over the public internet
B. Establish a Direct Connect to the VPC Route all traffic to the application via the direct connect connection
C. Enable VPC peering to decrease latency between instances Enable QoS on peering connection
D. Configure Amazon Trusted Advisor to give higher prioritization to the IP to assigned to the remote office over public internet traffic

Answer: C

Explanation:
https://docs.aws.amazon.com/vpc/latest/peering/create-vpc-peering-connection.html

NEW QUESTION 380
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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