Amazon-Web-Services

Exam Questions SOA-C01

AWS Certified SysOps Administrator - Associate

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

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

Answer: B

NEW QUESTION 2
When attached to an Amazon VPC which two components provide connectivity with external networks? Choose 2 answers

A. Elastic IPS (EIP)
B. NAT Gateway (NAT)
C. Internet Gateway {IGW)
D. Virtual Private Gateway (VGW)

Answer: CD

NEW QUESTION 3
Your application currently leverages AWS Auto Scaling to grow and shrink as load Increases/ decreases and has been performing well Your marketing team expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks Your forecast for the approximate number of Amazon EC2 instances necessary to meet the peak demand is 175.
What should you do to avoid potential service disruptions during the ramp up in traffic?

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

Answer: B

Explanation:
As the EC2 limit per region is max 20. You will need to fill an Amazon EC2 instance request form to increase the EC2 instances to 175.
http://aws.amazon.com/ec2/faqs/#How_many_instances_can_I_run_in_Amazon_EC2
I don??t think the answer can be D, as the question says ??expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks??. To pre-warm your ELB, you have to put in a request to AWS. You can??t do it.
Q: How do I reserve capacity for an existing, running instance?
To reserve capacity for a running instance, you can purchase a Reserved Instance or modify an existing reservation so it matches your instance's specifications. You can purchase Reserved Instances via the Amazon EC2 Console or by using the PurchaseReservedInstancesOffering API. You can modify existing Reserved Instances via the Amazon EC2 Console or by using the ModifyReservedInstances API call.
In both cases, the reservation must match the following attributes of the running instance you want to cover:
Availability Zone (e.g., us-east-1a) Instance type (e.g., m3.large)
Platform (e.g., Linux/UNIX (Amazon VPC)) Tenancy (e.g., default)
Q: How do I control which instances are billed at the lower rate?
The RunInstances API command does not distinguish between On-Demand instances and the reservations that can be applied to them. When computing your bill, our system will automatically optimize which instances are charged at the lower rate to ensure you always pay the lowest amount. For information about hourly billing, and how it applies to Reserved Instances, see Billing Benefits and Payment Options.
Q: How many Reserved Instances can I purchase?
You can purchase up to 20 Reserved Instances per Availability Zone each month. If you need additional Reserved Instances, complete the form found here. Information about previous generation Reserved Instance types can be found here.
Q: Can I reassign my Reserved Instance from one instance type (e.g., c1.xlarge) to another (e.g., m1.large)?
No. A Reserved Instance is associated with a specific instance type for the duration of its term; however, you can change from one instance size (e.g., c3.large) to another (e.g., c3.xlarge) in the same type, if it is a Linux/UNIX Reserved Instance.
Q: Can I move a Reserved Instance from one region to another?
No. A Reserved Instance is associated with a specific region, which is fixed for the duration of the reservation's term.
Q: Can I modify a Reserved Instance?
Yes. You can request to modify active reservations that you own in one of the following ways: Move between Availability Zones within the same region.
Change the network platform from EC2-Classic to EC2-VPC (for EC2-Classic-enabled customers). Change the instance type of your Linux/UNIX Reserved Instances to a larger or smaller size in the same instance type (e.g., convert 8 m1.smalls into 4 m1.mediums, or vice versa).
Instance type modifications are only supported for Linux/UNIX platform reservations. However, due to licensing differences Linux Reserved Instances cannot be modified to RedHat or SUSE Linux Reserved Instances.
The reservations that you modify must have been purchased on the same day, be the same instance type, and in the same Availability Zone and region. It is not possible to combine reservations. However, if you have multiple instances in the same reservation (i.e., the reservation was purchased to apply to 10 instances), you can modify each of these instances either individually or as a whole.
Q: How do I request changes or modifications?
You can submit a modification request from the Amazon EC2 Console or by using the ModifyReservedInstances API. We process your requests as soon as possible, depending on available capacity. There is no additional cost for modifying your Reserved Instances.
To learn more about modification, see the Amazon EC2 User Guide.

NEW QUESTION 4
The majority of your Infrastructure is on premises and you have a small footprint on AWS Your company has decided to roll out a new application that is heavily dependent on low latency connectivity to LOAP for authentication Your security policy requires minimal changes to the company's existing application user management processes.

What option would you implement to successfully launch this application1?

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

Answer: C

Explanation:
Since it requires no changes to the authentication infrastructure as requested in the question. Option D creates a new LDAP, trusts, etc.

NEW QUESTION 5
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 6
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 7
You have two Elastic Compute Cloud (EC2) instances inside a Virtual Private Cloud (VPC) in the same Availability Zone (AZ) but in different subnets. One instance is running a database and the other instance an application that will interface with the database. You want to confirm that they can talk to each other for your application to work properly.
Which two things do we need to confirm in the VPC settings so that these EC2 instances can communicate inside the VPC?
Choose 2 answers

A. A network ACL that allows communication between the two subnets.
B. Both instances are the same instance class and using the same Key-pair.
C. That the default route is set to a NAT instance or internet Gateway (IGW) for them to communicate.
D. Security groups are set to allow the application host to talk to the database on the right port/protocol.

Answer: AD

NEW QUESTION 8
You are managing a legacy application Inside VPC with hard coded IP addresses in its configuration. Which two mechanisms will allow the application to failover to new instances without the need for reconfiguration? Choose 2 answers

A. Create an ELB to reroute traffic to a failover instance
B. Create a secondary ENI that can be moved to a failover instance
C. Use Route53 health checks to fail traffic over to a failover instance
D. Assign a secondary private IP address to the primary ENI0 that can be moved to a failover instance

Answer: BD

Explanation:
This is an odd question. First of all, option A cannot be right because ELB does not failover. Cannot be C because Route 53 does work with hard coded IP. Only B

& D cannot be rule out so best answer.

NEW QUESTION 9
Which of the following statements about this S3 bucket policy is true?

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

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

Answer: B

NEW QUESTION 10
You have been asked to propose a multi-region deployment of a web-facing application where a controlled portion of your traffic is being processed by an alternate region.
Which configuration would achieve that goal?

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

Answer: A

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

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

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

Answer: A

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

NEW QUESTION 13
Your organization's security policy requires that all privileged users either use frequently rotated passwords or one-time access credentials in addition to username/password.
Which two of the following options would allow an organization to enforce this policy for AWS users? Choose 2 answers

A. Configure multi-factor authentication for privileged 1AM users
B. Create 1AM users for privileged accounts
C. Implement identity federation between your organization's Identity provider leveraging the 1AM Security Token Service
D. Enable the 1AM single-use password policy option for privileged users

Answer: AB

Explanation:
See also: http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html
Enable MFA for privileged users
For extra security, enable multifactor authentication (MFA) for privileged IAM users (users who are allowed access to sensitive resources or APIs). With MFA, users have a device that generates a unique authentication code (a one-time password, or OTP) and users must provide both their normal credentials (like their user name and password) and the OTP. The MFA device can either be a special piece of hardware, or it can be a virtual device (for example, it can run in an app on a smartphone).

NEW QUESTION 14
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 19
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 20
Your team is excited about the use of AWS because now they have access to ??programmable Infrastructure" You have been asked to manage your AWS infrastructure In a manner similar to the way you might manage application code You want to be able to deploy exact copies of different versions of your infrastructure, stage changes into different environments, revert back to previous versions, and identify what versions are running at any particular time (development, test, QA, production).
Which approach addresses this requirement?

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

Answer: D

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

NEW QUESTION 25
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 26
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 27
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 29
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 31
A user has deployed an application on his private cloud. The user is using his own monitoring tool. He wants to configure that whenever there is an error, the monitoring tool should notify him via SMS. Which of the below mentioned AWS services will help in this scenario?

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

Answer: B

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

NEW QUESTION 33
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 36
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 40
A user is trying to aggregate all the CloudWatch metric data of the last 1 week. Which of the below mentioned statistics is not available for the user as a part of data aggregation?

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

Answer: A

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

NEW QUESTION 41
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 46
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 47
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 51
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 55
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 60
A user is trying to configure the CloudWatch billing alarm. Which of the below mentioned steps should be performed by the user for the first time alarm creation in the AWS Account Management section?

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

Answer: B

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

NEW QUESTION 62

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 63
A user has launched an EBS backed EC2 instance. What will be the difference while performing the
restart or stop/start options on that instance?

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

Answer: A

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

NEW QUESTION 68
A sysadmin is trying to understand the Auto Scaling activities. Which of the below mentioned processes is not performed by Auto Scaling?

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

Answer: A

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

NEW QUESTION 73
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 76
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 78
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 81
A user is trying to setup a recurring Auto Scaling process. The user has setup one process to scale up every day at 8 am and scale down at 7 PM. The user is trying to setup another recurring process which scales up on the 1st of every month at 8 AM and scales down the same day at 7 PM. What will Auto Scaling do in this scenario?

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

Answer: D

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

NEW QUESTION 85
A user is trying to understand the ACL and policy for an S3 bucket. Which of the below mentioned policy permissions is equivalent to the WRITE ACL on a bucket?

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

Answer: D

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

NEW QUESTION 86
A user has launched an EC2 instance from an instance store backed AMI. The infrastructure team wants to create an AMI from the running instance. Which of the below mentioned steps will not be performed while creating the AMI?

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

Answer: A

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

NEW QUESTION 89
A user has created a VPC with CIDR 20.0.0.0/16. The user has created public and VPN only subnets along with hardware VPN access to connect to the user??s datacenter. The user wants to make so that all traffic coming to the public subnet follows the organization??s proxy policy. How can the user make this happen?

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

Answer: D

Explanation:
The user can create subnets within a VPC. If the user wants to connect to VPC from his own data centre, he can setup public and VPN only subnets which uses

hardware VPN access to connect with his data centre. When the user has configured this setup, it will update the main route table used with the VPN-only subnet, create a custom route table and associate it with the public subnet. It also creates an internet gateway for the public subnet. By default, the internet traffic of the VPN subnet is routed to a virtual private gateway while the internet traffic of the public subnet is routed through the internet gateway. The user can set up the route and security group rules. These rules enable the traffic to come from the organization??s network over the virtual private gateway to the public subnet to allow proxy settings on that public subnet.

NEW QUESTION 92
A user is using a small MySQL RDS DB. The user is experiencing high latency due to the Multi AZ feature. Which of the below mentioned options may not help the user in this situation?

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

Answer: D

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

NEW QUESTION 97
A user is planning to schedule a backup for an EBS volume. The user wants security of the snapshot data. How can the user achieve data encryption with a snapshot?

A. Use encrypted EBS volumes so that the snapshot will be encrypted by AWS
B. While creating a snapshot select the snapshot with encryption
C. By default the snapshot is encrypted by AWS
D. Enable server side encryption for the snapshot using S3

Answer: A

Explanation:
AWS EBS supports encryption of the volume. It also supports creating volumes from existing snapshots provided the snapshots are created from encrypted volumes. The data at rest, the I/O as well as all the snapshots of the encrypted EBS will also be encrypted. EBS encryption is based on the AES-256 cryptographic algorithm, which IT the industry standard.

NEW QUESTION 101
A user has setup an EBS backed instance and attached 2 EBS volumes to it. The user has setup a CloudWatch alarm on each volume for the disk data. The user has stopped the EC2 instance and detached the EBS volumes. What will be the status of the alarms on the EBS volume?

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

Answer: B

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

NEW QUESTION 103
A user has launched an EC2 instance from an instance store backed AMI. The infrastructure team wants to create an AMI from the running instance. Which of the below mentioned credentials is not required while creating the AMI?

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

Answer: C

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

NEW QUESTION 104
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 107
An organization is measuring the latency of an application every minute and storing data inside a file in the JSON format. The organization wants to send all latency data to AWS CloudWatch. How can the organization achieve this?

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

Answer: C

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

NEW QUESTION 108
A user has launched an EBS backed instance with EC2-Classic. The user stops and starts the instance. Which of the below mentioned statements is not true with respect to the stop/start action?

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

Answer: C

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

NEW QUESTION 113
A user has launched multiple EC2 instances for the purpose of development and testing in the same region. The user wants to find the separate cost for the production and development instances. How can the user find the cost distribution?

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

Answer: D

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

NEW QUESTION 115
A user has created a VPC with CIDR 20.0.0.0/16 using VPC Wizard. The user has created a public CIDR (20.0.0.0/24. and a VPN only subnet CIDR (20.0.1.0/24. along with the hardware VPN access to connect to the user??s data centre. Which of the below mentioned components is not present when the VPC is setup with the wizard?

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

Answer: B

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

NEW QUESTION 116
A user has setup an Auto Scaling group. The group has failed to launch a single instance for more than 24 hours. What will happen to Auto Scaling in this condition?

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

Answer: B

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

NEW QUESTION 119
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 123
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 124
A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered. What time out period should the user specify for connection draining?

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

Answer: B

NEW QUESTION 127
A user is using the AWS EC2. The user wants to make so that when there is an issue in the EC2 server, such as instance status failed, it should start a new instance in the user??s private cloud. Which AWS service helps to achieve this automation?

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

D. AWS CloudWatch + AWS SNS

Answer: D

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

NEW QUESTION 128
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 131
A user has enabled session stickiness with ELB. The user does not want ELB to manage the cookie; instead he wants the application to manage the cookie. What will happen when the server instance, which is bound to a cookie, crashes?

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

Answer: B

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

NEW QUESTION 136
A user has created a VPC with public and private subnets using the VPC wizard. Which of the below mentioned statements is not true in this scenario?

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

Answer: A

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

NEW QUESTION 137
A user has configured ELB with a TCP listener at ELB as well as on the back-end instances. The user wants to enable a proxy protocol to capture the source and destination IP information in the header. Which of the below mentioned statements helps the user understand a proxy protocol with TCP configuration?

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

Answer: A

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

NEW QUESTION 140
A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. Which of the below mentioned statements is true with respect to this scenario?

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

Answer: D

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

NEW QUESTION 144
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 145
A user had aggregated the CloudWatch metric data on the AMI ID. The user observed some abnormal behaviour of the CPU utilization metric while viewing the last 2 weeks of data. The user wants to share that data with his manager. How can the user achieve this easily with the AWS console?

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

Answer: A

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

NEW QUESTION 148
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 151
A user has configured ELB with SSL using a security policy for secure negotiation between the client and load balancer. Which of the below mentioned SSL protocols is not supported by the security policy?

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

Answer: A

Explanation:
Elastic Load Balancing uses a Secure Socket Layer (SSL. negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. Elastic Load Balancing supports the following versions of the SSL protocol:
TLS 1.2
TLS 1.1
TLS 1.0
SSL 3.0
SSL 2.0

NEW QUESTION 153
A user has launched an EC2 Windows instance from an instance store backed AMI. The user has also set the Instance initiated shutdown behavior to stop. What will happen when the user shuts down the OS?

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

Answer: B

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

NEW QUESTION 157
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 159
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 162
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 167
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 170
A user has created a VPC with public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24 . The NAT instance ID is i-a12345. Which of the below mentioned entries are required in the main route table attached with the private subnet to allow instances to connect with the internet?

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

Answer: A

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

NEW QUESTION 171
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 173
An AWS account owner has setup multiple IAM users. One IAM user only has CloudWatch access. He has setup the alarm action which stops the EC2 instances when the CPU utilization is below the threshold limit. What will happen in this case?

A. It is not possible to stop the instance using the CloudWatch alarm
B. CloudWatch will stop the instance when the action is executed
C. The user cannot set an alarm on EC2 since he does not have the permission

D. The user can setup the action but it will not be executed if the user does not have EC2 rights

Answer: D

Explanation:
Amazon CloudWatch alarms watch a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup an action which stops 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. If the IAM user has read/write permissions for Amazon CloudWatch but not for Amazon EC2, he can still create an alarm. However, the stop or terminate actions will not be performed on the Amazon EC2 instance.

NEW QUESTION 175
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 by mistake. The user is trying to create another subnet of CIDR 20.0.0.1/24. How can the user create the second subnet?

A. There is no need to update the subnet as VPC automatically adjusts the CIDR of the first subnet based on the second subnet??s CIDR
B. The user can modify the first subnet CIDR from the console
C. It is not possible to create a second subnet as one subnet with the same CIDR as the VPC has been created
D. The user can modify the first subnet CIDR with AWS CLI

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 the 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. The user cannot modify the CIDR of a subnet once it is created. Thus, in this case if required, the user has to delete the subnet and create new subnets.

NEW QUESTION 176
A user has created a VPC with the public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The public subnet uses CIDR 20.0.1.0/24. The user is planning to host a web server in the public subnet (port 80. and a DB server in the private subnet (port 3306.. The user is configuring a security group for the public subnet (WebSecGrp. and the private subnet (DBSecGrp.. Which of the below mentioned entries is required in the web server security group (WebSecGrp.?

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

Answer: A

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

NEW QUESTION 177
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 182
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 187
A user has granted read/write permission of his S3 bucket using ACL. Which of the below mentioned options is a valid ID to grant permission to other AWS accounts (grantee. using ACL?

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

Answer: D

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

NEW QUESTION 188
A user has created an Auto Scaling group using CLI. The user wants to enable CloudWatch detailed monitoring for that group. How can the user configure this?

A. When the user sets an alarm on the Auto Scaling group, it automatically enables detail monitoring
B. By default detailed monitoring is enabled for Auto Scaling
C. Auto Scaling does not support detailed monitoring
D. Enable detail monitoring from the AWS console

Answer: B

Explanation:
CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. To enable detailed instance monitoring for a new Auto Scaling group, the user does not need to take any extra steps. When the user creates an Auto Scaling launch config as the first step for creating an Auto Scaling group, each launch configuration contains a flag named InstanceMonitoring.Enabled. The default value of this flag is true. Thus, the user does not need to set this flag if he wants detailed monitoring.

NEW QUESTION 189
A user has created a VPC with a public subnet. The user has terminated all the instances which are part of the subnet. Which of the below mentioned statements is true with respect to this scenario?

A. The user cannot delete the VPC since the subnet is not deleted
B. All network interface attached with the instances will be deleted
C. When the user launches a new instance it cannot use the same subnet
D. The subnet to which the instances were launched with will be deleted

Answer: B

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

NEW QUESTION 190
A user has created an application which will be hosted on EC2. The application makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK to connect with from the EC2 instance. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

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

Answer: A

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

NEW QUESTION 193
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 194
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 198
A user is trying to create a PIOPS EBS volume with 8 GB size and 200 IOPS. Will AWS create the volume?

A. Yes, since the ratio between EBS and IOPS is less than 30
B. No, since the PIOPS and EBS size ratio is less than 30
C. No, the EBS size is less than 10 GB
D. Yes, since PIOPS is higher than 100

Answer: C

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

NEW QUESTION 203
An organization has launched 5 instances: 2 for production and 3 for testing. The organization wants that one particular group of IAM users should only access the test instances and not the production ones. How can the organization set that as a part of the policy?

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

Answer: D

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

```
}
}
}
]
```

NEW QUESTION 206
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 211
A user wants to upload a complete folder to AWS S3 using the S3 Management console. How can the user perform this activity?

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

Answer: D

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

NEW QUESTION 215
Which of the below mentioned AWS RDS logs cannot be viewed from the console for MySQL?

A. Error Log
B. Slow Query Log
C. Transaction Log
D. General Log

Answer: C

Explanation:
The user can view, download, and watch the database logs using the Amazon RDS console, the Command Line Interface (CLI., or the Amazon RDS API. For the MySQL RDS, the user can view the error log, slow querylog, and general logs. RDS does not support viewing the transaction logs.

NEW QUESTION 220
An organization has created 10 IAM users. The organization wants each of the IAM users to have access to a separate DyanmoDB table. All the users are added to the same group and the organization wants to setup a group level policy for this. How can the organization achieve this?

A. Define the group policy and add a condition which allows the access based on the IAM name
B. Create a DynamoDB table with the same name as the IAM user name and define the policy rule which grants access based on the DynamoDB ARN using a variable
C. Create a separate DynamoDB database for each user and configure a policy in the group based on the DB variable
D. It is not possible to have a group level policy which allows different IAM users to different DynamoDB Tables

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. AWS DynamoDB has only tables and the organization cannot makeseparate databases. The organization should create a table with the same name as the IAM user name and use the ARN of DynamoDB as part of the group policy. The sample policy is shown below:
```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": ["dynamodb:*"],
"Resource": "arn:aws:dynamodb:region:account-number-without-hyphens:table/${aws:username}"
}
]
}
```

NEW QUESTION 224
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 229
A sys admin has enabled a log on ELB. Which of the below mentioned activities are not captured by the log?

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

Answer: B

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

NEW QUESTION 234
A user runs the command ??dd if=/dev/xvdf of=/dev/null bs=1M?? on an EBS volume created from a snapshot and attached to a Linux instance. Which of the below mentioned activities is the user performing with the step given above?

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

Answer: A

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

NEW QUESTION 235
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 240
Which method can be used to prevent an IP address block from accessing public objects in an S3 bucket?

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

Answer: A

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

NEW QUESTION 242
What would happen to an RDS (Relational Database Service) multi-Availability Zone deployment if the primary DB instance fails?

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

Answer: D

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

NEW QUESTION 247
How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

A. Simply create a new volume in the other AZ and specify the original volume as the source.
B. Detach the volume, then use the ec2-migrate-volume command to move it to another AZ.
C. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
D. Detach the volume and attach it to another EC2 instance in the other AZ.

Answer: C

Explanation:
Snapshots can be used to create multiple new EBS volumes, expand the size of a volume, or move volumes across Availability Zones.
See: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumes.html

NEW QUESTION 249
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 250
A company needs to monitor the read and write IOPs metrics for their AWS MySQL RDS instance and send real-time alerts to their operations team. Which AWS services can accomplish this? Choose 2 answers

A. Amazon Simple Email Service
B. Amazon CloudWatch
C. Amazon Simple Queue Service
D. Amazon Route 53
E. Amazon Simple Notification Service

Answer: BE

NEW QUESTION 254
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 255
Which features can be used to restrict access to data in S3? Choose 2 answers

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

Answer: AC

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

NEW QUESTION 259
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 263
A syslog Administrator is created additional Amazon EC2 instances and receive an Instancelimitexceeded error.
What is the cause of the issue and how can it be resolve?

A. The Administrator has requested too many instances at once and must request fewer instances in batches
B. The concurrent running instance limit has been reached and an EC2 limit increase request must be filed with AWS Support
C. AWS does not currently nave enough available capacity and a different instance type must be used
D. The Administrator must specify the maximum number of instances to be ?V created provisioning EC stances

Answer: B

Explanation:
EC2 Service Limits: AWS sets limits for these resources on a per-region basis.
If you are getting an InstanceLimitExceeded error when you try to launch an instance, you have reached your concurrent running instance limit. For new AWS accounts, the default limit is 20. If you need additional running instances, complete the form at Request to Increase Amazon EC2 Instance Limit.
By default, all AWS accounts have a limit of 20 running instances at any time per region. If you attempt to start another one, even if it already existed in the stopped state, you will receive this error message.
To resolve this issue, you can do any of the following: Stop one of your other running instances
Contact AWS support and request your running EC2 instances quota limit be raised.

NEW QUESTION 264
A Developers that an Amazon EC2 instance has failed. The developer reports that all the data was stored on the root volume is now gone.
What is the expiation for this issue?

A. The instance was using an Amazon EBS root volume
B. The instance was using Amazon S3 as the root volume
C. The instance was using an instance store root volume
D. The root volume with the data exists but needs to be re-attached

Answer: A

Explanation:
If your instance is ebs-backed, then you will not lose ebs root volume storage if you launched it with the "delete-on-termination" set to false. See the ec2-run-instances command for more information. Other ephemeral volumes will be lost when the instance is stopped/terminated.
If it is not ebs-backed, you will lose the root data when you terminate the instance (you cannot "stop" instance store instances).

NEW QUESTION 268
The Security tram is connect because the number of AWS identity and access Management (IAM) policies being in the environment is increasing. The tasked a SysOps Administrator to report on the number of IAM policies in use and use the total IAM policies.
Which AWS service should the Administrator use to check how current IAM policy compares to current limits?

A. MWS Trusted Advisor
B. Amazon Inspector
C. AWS Config
D. Organizations

Answer: C

Explanation:
AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This enables you to simplify compliance auditing, security analysis, change management, and operational troubleshooting.

NEW QUESTION 271
A company has created a separate AWS account for all development work to protect the production environment in this development account, developers have permission to manipulate IAM policies and roles. Corporate policies require that developers and blocked from accessing some services. What is the BEST way to grant the developers privileges in the development account while still complying with corporate policies?

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

Answer: A

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

NEW QUESTION 272
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 276
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 281
A company is migrating an application to AWS that requires access to a legacy system, which remain in the company's data centre. The application runs inside a VPC in the company's AWS account. The application must offer a consistent and low-latency response to its users
How can these requirements be met?

A. Create a software-based VPN connection between the Amazon VPC and the on-premises network
B. Create an AWS Direct Connect connection between AWS and the on-premises network and then use a private virtual interface

C. Create a hardware-based IPsec VPN connection between the VPC in AWS and the on-premises network
D. Create an overlay network by using third-party software and use that to connect the X/PC back to the on-premises network

Answer: B

Explanation:
Private Connectivity to your Amazon VPC. You can use AWS Direct Connect to establish a private virtual interface from your on-premise network directly to your Amazon VPC, providing you with a private, high bandwidth network connection between your network and your VPC.

NEW QUESTION 284
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 289
A SysOps Administrator has set up a new Application Load Balancer (ALB) in front of a pair of private web server in multiple Availability Zones. After deployment an updates CloudFormation template with many changes, user now goes to one web server only.
What is the NOST likely reason that the traffic is not being balanced between both servers?

A. The faulty is returning HTTP 200 has been removed.
B. Sticky session have been disabled in the ALB for the working sever.
C. The ALB using a custom ping path that is not found on the faulty server.
D. The web client are using HTTP/2, which is terminated at the ALB.

Answer: B

Explanation:
Until now, the behavior of load balancers has been to route each request independently to the Amazon EC2 instance with the least load. With the stickiness feature, you can configure the load balancer to bind user sessions to specific application instances. All requests coming from the user during the session will be sent to the same application instance. Elastic Load Balancing supports two mechanisms to provide session stickiness: load balancer-generated HTTP cookies, which allow browser-based session lifetimes, and application-generated HTTP cookies, which allow application- specific session lifetimes. You can learn more about this feature by visiting the ELB Developers Guide.

NEW QUESTION 290
A System Administrator is trying to identify why Put Object calls are not made from an Amazon EC2 instance to an Amazon bucket in the same region.
The instance is launched in a subnet with CIDR range 10.1.0.24 and 'Auto assign public IP set to yes. The instance profile tied to this instance has AmazonS3Access policy.
Security group rules for the instance:

Protocol	PortRange	Source
HTTP (80)	80	0.0.0.0/0
HTTPS (443)	443	0.0.0.0/0
Custom TCP	1024-65535	0.0.0.0/0

The route table for the subnet in which this instance is launched

Destination	Target
10.0.0.0/16	local

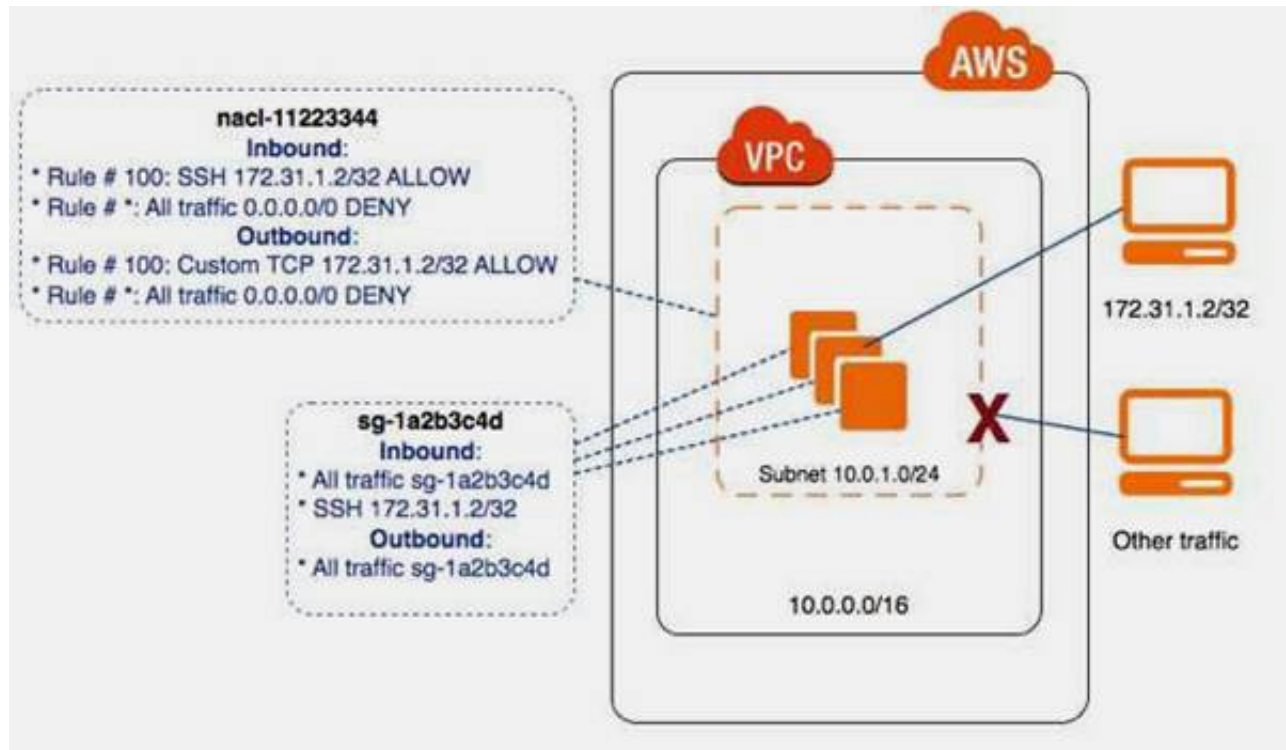Based on the information provided what is causing the lack of access to S3 from the instance?

A. The instance profile does not have explicit permissions to write objects to the S3 bucket.
B. The route table does not have a rule for all traffic to pass through a NAT gateway.
C. The route table does not have rule for all traffic to pass through an internet gateway

Answer: B

Explanation:
Controlling Access to Instances in a Subnet
In this example, instances in your subnet can communicate with each other, and are accessible from a trusted remote computer. The remote computer may be a computer in your local network or an instance in a different subnet or VPC that you use to connect to your instances to perform administrative tasks. Your security group rules and network ACL rules allow access from the IP address of your remote computer (172.31.1.2/32). All other traffic from the Internet or other networks is denied.

All instances use the same security group (sg-1a2b3c4d), with the following rules.
Protocol Protocol Port Source Comments

Type		Range		
All traffic	All	All	sg-1a2b3c4d	Enables instances associated with the same security group to communicate with each other.
TCP	SSH	22	172.31.1.2/32	Allows inbound SSH access from the remote computer. If the instance is a Windows computer, then this rule must use the RDP protocol for port 3389 instead.

Protocol Type	Protocol	Port Range	Destination	Comments
All traffic	All	All	sg-1a2b3c4d	Enables instances associated with the same security group to communicate with each other.

The subnet is associated with a network ACL that has the following rules.

Rule #	Type	Protocol	Port Range	Source	Allow/Deny	Comments
100	SSH	TCP	22	172.31.1.2/32	ALLOW	Allows inbound traffic from the remote computer. If the instance is a Windows computer, then this rule must use the RDP protocol for port 3389 instead.
*	All traffic All	All	All	0.0.0.0/0	DENY	Denies all other inbound traffic that does not match the previous rule.

Rule #	Type	Protocol	Port Range	Destination	Allow/Deny	Comments
100	Custom TCP	TCP	1024-65535	172.31.1.2/32	ALLOW	Allows outbound responses to the remote computer. Network ACLs are stateless, therefore this rule is required to allow response traffic for inbound requests.
*	All traffic All	All	All	0.0.0.0/0	DENY	Denies all other outbound traffic that does not match the previous rule.

This scenario gives you the flexibility to change the security groups or security group rules for your instances, and have the network ACL as the backup layer of defense. The network ACL rules apply to all instances in the subnet, so if you accidentally make your security group rules too permissive, the network ACL rules continue to permit access only from the single IP address. For example, the following rules are more permissive than the earlier rules ?X they allow inbound SSH access from any IP address.

Type	Protocol	Port Range	Source	Comments
All traffic	All	All	sg-1a2b3c4d	Enables instances associated with the same security group to communicate with each other.
SSH	TCP	22	0.0.0.0/0	Allows SSH access from any IP address.

Type	Protocol	Port Range	Destination	Comments
All traffic	All	All	0.0.0.0/0	Allows all outbound traffic.

However, only other instances within the subnet and your remote computer are able to access this instance. The network ACL rules still prevent all inbound traffic to the subnet except from your remote computer.

NEW QUESTION 292
An Application team is using Remote Desktop to connect to its application server and perform administrative tasks. After deployment a Windows service a existing subnets, the team discovers that it is unable to communicate with the new servers. A SysOps Administrative has obtained the VPC logs as shown in the table) related to the communication to help troubleshooting the problem.

How can this issue be resolved?

A. Check the route Tables to validate that the Remote Desktop and return traffic is allowed to and from the new servers.
B. Check the security groups to validate that Remote Desktop is allowed into the new servers.
C. Check the network access control lists to validate that the Remote Desktop and return traffic is allowed to and from the new servers.
D. Ensures that the RDP service and Windows firewall are open and listening on Port 3389 TCP.

Answer: D

NEW QUESTION 297
A company operate a secure website running an Amazon EC2 instance behind a Classic Load Balancer. An SSL certificate from AWS Certificate Manager is deployment on the load balancer. The company's Marketing team has determined that too many customer using older browser are experiencing issues with the website has asked a SysOps Administrator to fix this issue.
What course of action should the administrator take?

A. Update the SSL negotiation configuration of the load balancer by creating a custom security polic
B. Ensure the appropriate cipher has been enabled so that the web application can support the webbrowser.
C. Create a separate Classic Load Balancer and install custom SSL certificate with a different domain name on it that support the web browse
D. Ask customer with the affected browser to use this domain name instead of the one they are accustomed to using.
E. Create a new SSL certificate in Certificate Manager and install this certificate on each of the servers to accommodates the web browsers.
F. Remove the load balancer from the configuration and instead install a custom SSL certificate on each of the web servers.

Answer: A

Explanation:
Update the SSL Negotiation Configuration of Your Classic Load Balancer
Elastic Load Balancing provides security policies that have predefined SSL negotiation configurations to use to negotiate SSL connections between clients and your load balancer. If you are using the HTTPS/SSL protocol for your listener, you can use one of the predefined security policies, or use your own custom security policy.
For more information about the security policies, see SSL Negotiation Configurations for Classic Load Balancers. For information about the configurations of the security policies provided by Elastic Load Balancing, see Predefined SSL Security Policies.
If you create an HTTPS/SSL listener without associating a security policy, Elastic Load Balancing associates the default predefined security policy, ELBSecurityPolicy-2016-08, with your load balancer. If you have an existing load balancer with an SSL negotiation configuration that does not use the latest protocols and ciphers, we recommend that you update your load balancer to use ELBSecurityPolicy-2016-08. If you prefer, you can create a custom configuration. We strongly recommend that you test the new security policies before you upgrade your load balancer configuration.
The following examples show you how to update the SSL negotiation configuration for an HTTPS/SSL listener. Note that the change does not affect requests that were received by a load balancer node and are pending routing to a healthy instance, but the updated configuration will be used with new requests that are received.

NEW QUESTION 301
A company has mandated the use factor authentication (MFA) for all user, and requires users to make all API calls using CLI. However, uses are not prompted to enter MFA token, and able to return CLI commands without MF

A. In an enforce MFA, the company attached an IAM policy to all users that derives API calls that not been authenticated with MF
B. What additional step must be ensure that calls are authenticated using MFA?
C. Enable MFA on IAM roles, requires IAM to use role credentials to sign API calls.
D. Ask the IAM to log into the AWS Management Console with MFA before marking PI calls using the Cli.
E. Restricted the IAM users to use the console, as MFA not supported for CLI use.
F. Reporting users to use temporary credential from the get-session token command to sign API calls.

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/authenticate-mfa-cli/

NEW QUESTION 302
A SysOps Administrator has an AWS Lambda function that performs maintenance on versions AWS resources. This function must be run nightly.
Which is the MOST cost-effective solution?

A. Launch a single l2.nano Amazon EC2 instance and create a Linux corn job to invoke the Lambda function at the same every right.
B. Set up an Amazon CloudWatch metric alarm to invoke the Lambda function at the same time every night.
C. Schedule a CloudWatch event to invoke the Lambda function at the same time every night.
D. Implement a Chef recipe in Opsworks stack to invoke the Lambda function at the same time every night

Answer: C

Explanation:
Using AWS Lambda with Amazon CloudWatch Events
You can create a Lambda function and direct AWS Lambda to execute it on a regular schedule. You can specify a fixed rate (for example, execute a Lambda function every hour or 15 minutes), or you can specify a Cron expression. For more information on expressions schedules, see Schedule Expressions Using Rate or Cron.
This functionality is available when you create a Lambda function using the AWS Lambda console or the AWS CLI. To configure it using the AWS CLI, see Run an AWS Lambda Function on a Schedule Using the AWS CLI. The console provides CloudWatch Events as an event source. At the time of creating a Lambda function, you choose this event source and specify a time interval.
If you have made any manual changes to the permissions on your function, you may need to reapply
the scheduled event access to your function. You can do that by using the following CLI command.
$ aws lambda add-permission --function-name function_name\
--action 'lambda:InvokeFunction' --principal events.amazonaws.com \
--statement-id 'statement_id' \
--source-arn arn:aws:events:region:account-id:rule/rule_name

Each AWS account can have up to 100 unique event sources of the CloudWatch Events- Schedule source type. Each of these can be the event source for up to five Lambda functions. That is, you can have up to 500 Lambda functions that can be executing on a schedule in your AWS account.

The console also provides a blueprint (lambda-canary) that uses the CloudWatch Events - Schedule source type. Using this blueprint, you can create a sample Lambda function and test this feature. The example code that the blueprint provides checks for the presence of a specific webpage and specific text string on the webpage. If either the webpage or the text string is not found, the Lambda function throws an error.

NEW QUESTION 306

......

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

    All examinations will be up to date.

* 24/7 Quality Support

    We will provide service round the clock.

* 100% Pass Rate

    Our guarantee that you will pass the exam.

* Unique Gurantee

    If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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

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

Answer: B

NEW QUESTION 2
When attached to an Amazon VPC which two components provide connectivity with external networks? Choose 2 answers

A. Elastic IPS (EIP)
B. NAT Gateway (NAT)
C. Internet Gateway {IGW)
D. Virtual Private Gateway (VGW)

Answer: CD

NEW QUESTION 3
Your application currently leverages AWS Auto Scaling to grow and shrink as load Increases/ decreases and has been performing well Your marketing team expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks Your forecast for the approximate number of Amazon EC2 instances necessary to meet the peak demand is 175.
What should you do to avoid potential service disruptions during the ramp up in traffic?

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

Answer: B

Explanation:
As the EC2 limit per region is max 20. You will need to fill an Amazon EC2 instance request form to increase the EC2 instances to 175.
http://aws.amazon.com/ec2/faqs/#How_many_instances_can_I_run_in_Amazon_EC2
I don??t think the answer can be D, as the question says ??expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks??. To pre-warm your ELB, you have to put in a request to AWS. You can??t do it.
Q: How do I reserve capacity for an existing, running instance?
To reserve capacity for a running instance, you can purchase a Reserved Instance or modify an existing reservation so it matches your instance's specifications. You can purchase Reserved Instances via the Amazon EC2 Console or by using the PurchaseReservedInstancesOffering API. You can modify existing Reserved Instances via the Amazon EC2 Console or by using the ModifyReservedInstances API call.
In both cases, the reservation must match the following attributes of the running instance you want to cover:
Availability Zone (e.g., us-east-1a) Instance type (e.g., m3.large)
Platform (e.g., Linux/UNIX (Amazon VPC)) Tenancy (e.g., default)
Q: How do I control which instances are billed at the lower rate?
The RunInstances API command does not distinguish between On-Demand instances and the reservations that can be applied to them. When computing your bill, our system will automatically optimize which instances are charged at the lower rate to ensure you always pay the lowest amount. For information about hourly billing, and how it applies to Reserved Instances, see Billing Benefits and Payment Options.
Q: How many Reserved Instances can I purchase?
You can purchase up to 20 Reserved Instances per Availability Zone each month. If you need additional Reserved Instances, complete the form found here. Information about previous generation Reserved Instance types can be found here.
Q: Can I reassign my Reserved Instance from one instance type (e.g., c1.xlarge) to another (e.g., m1.large)?
No. A Reserved Instance is associated with a specific instance type for the duration of its term; however, you can change from one instance size (e.g., c3.large) to another (e.g., c3.xlarge) in the same type, if it is a Linux/UNIX Reserved Instance.
Q: Can I move a Reserved Instance from one region to another?
No. A Reserved Instance is associated with a specific region, which is fixed for the duration of the reservation's term.
Q: Can I modify a Reserved Instance?
Yes. You can request to modify active reservations that you own in one of the following ways: Move between Availability Zones within the same region.
Change the network platform from EC2-Classic to EC2-VPC (for EC2-Classic-enabled customers). Change the instance type of your Linux/UNIX Reserved Instances to a larger or smaller size in the same instance type (e.g., convert 8 m1.smalls into 4 m1.mediums, or vice versa).
Instance type modifications are only supported for Linux/UNIX platform reservations. However, due to licensing differences Linux Reserved Instances cannot be modified to RedHat or SUSE Linux Reserved Instances.
The reservations that you modify must have been purchased on the same day, be the same instance type, and in the same Availability Zone and region. It is not possible to combine reservations. However, if you have multiple instances in the same reservation (i.e., the reservation was purchased to apply to 10 instances), you can modify each of these instances either individually or as a whole.
Q: How do I request changes or modifications?
You can submit a modification request from the Amazon EC2 Console or by using the ModifyReservedInstances API. We process your requests as soon as possible, depending on available capacity. There is no additional cost for modifying your Reserved Instances.
To learn more about modification, see the Amazon EC2 User Guide.

NEW QUESTION 4
The majority of your Infrastructure is on premises and you have a small footprint on AWS Your company has decided to roll out a new application that is heavily dependent on low latency connectivity to LOAP for authentication Your security policy requires minimal changes to the company's existing application user management processes.

What option would you implement to successfully launch this application1?

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

Answer: C

Explanation:
Since it requires no changes to the authentication infrastructure as requested in the question. Option D creates a new LDAP, trusts, etc.

NEW QUESTION 5
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 6
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 7
You have two Elastic Compute Cloud (EC2) instances inside a Virtual Private Cloud (VPC) in the same Availability Zone (AZ) but in different subnets. One instance is running a database and the other instance an application that will interface with the database. You want to confirm that they can talk to each other for your application to work properly.
Which two things do we need to confirm in the VPC settings so that these EC2 instances can communicate inside the VPC?
Choose 2 answers

A. A network ACL that allows communication between the two subnets.
B. Both instances are the same instance class and using the same Key-pair.
C. That the default route is set to a NAT instance or internet Gateway (IGW) for them to communicate.
D. Security groups are set to allow the application host to talk to the database on the right port/protocol.

Answer: AD

NEW QUESTION 8
You are managing a legacy application Inside VPC with hard coded IP addresses in its configuration. Which two mechanisms will allow the application to failover to new instances without the need for reconfiguration? Choose 2 answers

A. Create an ELB to reroute traffic to a failover instance
B. Create a secondary ENI that can be moved to a failover instance
C. Use Route53 health checks to fail traffic over to a failover instance
D. Assign a secondary private IP address to the primary ENI0 that can be moved to a failover instance

Answer: BD

Explanation:
This is an odd question. First of all, option A cannot be right because ELB does not failover. Cannot be C because Route 53 does work with hard coded IP. Only B

& D cannot be rule out so best answer.

NEW QUESTION 9
Which of the following statements about this S3 bucket policy is true?

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

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

Answer: B

NEW QUESTION 10
You have been asked to propose a multi-region deployment of a web-facing application where a controlled portion of your traffic is being processed by an alternate region.
Which configuration would achieve that goal?

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

Answer: A

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

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

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

Answer: A

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

NEW QUESTION 13
Your organization's security policy requires that all privileged users either use frequently rotated passwords or one-time access credentials in addition to username/password.
Which two of the following options would allow an organization to enforce this policy for AWS users? Choose 2 answers

A. Configure multi-factor authentication for privileged 1AM users
B. Create 1AM users for privileged accounts
C. Implement identity federation between your organization's Identity provider leveraging the 1AM Security Token Service
D. Enable the 1AM single-use password policy option for privileged users

Answer: AB

Explanation:
See also: http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html
Enable MFA for privileged users
For extra security, enable multifactor authentication (MFA) for privileged IAM users (users who are allowed access to sensitive resources or APIs). With MFA, users have a device that generates a unique authentication code (a one-time password, or OTP) and users must provide both their normal credentials (like their user name and password) and the OTP. The MFA device can either be a special piece of hardware, or it can be a virtual device (for example, it can run in an app on a smartphone).

NEW QUESTION 14
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 19
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 20
Your team is excited about the use of AWS because now they have access to ??programmable Infrastructure" You have been asked to manage your AWS infrastructure In a manner similar to the way you might manage application code You want to be able to deploy exact copies of different versions of your infrastructure, stage changes into different environments, revert back to previous versions, and identify what versions are running at any particular time (development, test, QA, production).
Which approach addresses this requirement?

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

Answer: D

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

NEW QUESTION 25
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 26
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 27
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 29
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 31
A user has deployed an application on his private cloud. The user is using his own monitoring tool. He wants to configure that whenever there is an error, the monitoring tool should notify him via SMS. Which of the below mentioned AWS services will help in this scenario?

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

Answer: B

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

NEW QUESTION 33
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 36
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 40
A user is trying to aggregate all the CloudWatch metric data of the last 1 week. Which of the below mentioned statistics is not available for the user as a part of data aggregation?

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

Answer: A

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

NEW QUESTION 41
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 46
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 47
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 51
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 55
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 60
A user is trying to configure the CloudWatch billing alarm. Which of the below mentioned steps should be performed by the user for the first time alarm creation in the AWS Account Management section?

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

Answer: B

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

NEW QUESTION 62

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 63
A user has launched an EBS backed EC2 instance. What will be the difference while performing the
restart or stop/start options on that instance?

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

Answer: A

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

NEW QUESTION 68
A sysadmin is trying to understand the Auto Scaling activities. Which of the below mentioned processes is not performed by Auto Scaling?

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

Answer: A

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

NEW QUESTION 73
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 76
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 78
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 81
A user is trying to setup a recurring Auto Scaling process. The user has setup one process to scale up every day at 8 am and scale down at 7 PM. The user is trying to setup another recurring process which scales up on the 1st of every month at 8 AM and scales down the same day at 7 PM. What will Auto Scaling do in this scenario?

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

Answer: D

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

NEW QUESTION 85
A user is trying to understand the ACL and policy for an S3 bucket. Which of the below mentioned policy permissions is equivalent to the WRITE ACL on a bucket?

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

Answer: D

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

NEW QUESTION 86
A user has launched an EC2 instance from an instance store backed AMI. The infrastructure team wants to create an AMI from the running instance. Which of the below mentioned steps will not be performed while creating the AMI?

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

Answer: A

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

NEW QUESTION 89
A user has created a VPC with CIDR 20.0.0.0/16. The user has created public and VPN only subnets along with hardware VPN access to connect to the user??s datacenter. The user wants to make so that all traffic coming to the public subnet follows the organization??s proxy policy. How can the user make this happen?

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

Answer: D

Explanation:
The user can create subnets within a VPC. If the user wants to connect to VPC from his own data centre, he can setup public and VPN only subnets which uses

hardware VPN access to connect with his data centre. When the user has configured this setup, it will update the main route table used with the VPN-only subnet, create a custom route table and associate it with the public subnet. It also creates an internet gateway for the public subnet. By default, the internet traffic of the VPN subnet is routed to a virtual private gateway while the internet traffic of the public subnet is routed through the internet gateway. The user can set up the route and security group rules. These rules enable the traffic to come from the organization??s network over the virtual private gateway to the public subnet to allow proxy settings on that public subnet.

NEW QUESTION 92
A user is using a small MySQL RDS DB. The user is experiencing high latency due to the Multi AZ feature. Which of the below mentioned options may not help the user in this situation?

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

Answer: D

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

NEW QUESTION 97
A user is planning to schedule a backup for an EBS volume. The user wants security of the snapshot data. How can the user achieve data encryption with a snapshot?

A. Use encrypted EBS volumes so that the snapshot will be encrypted by AWS
B. While creating a snapshot select the snapshot with encryption
C. By default the snapshot is encrypted by AWS
D. Enable server side encryption for the snapshot using S3

Answer: A

Explanation:
AWS EBS supports encryption of the volume. It also supports creating volumes from existing snapshots provided the snapshots are created from encrypted volumes. The data at rest, the I/O as well as all the snapshots of the encrypted EBS will also be encrypted. EBS encryption is based on the AES-256 cryptographic algorithm, which IT the industry standard.

NEW QUESTION 101
A user has setup an EBS backed instance and attached 2 EBS volumes to it. The user has setup a CloudWatch alarm on each volume for the disk data. The user has stopped the EC2 instance and detached the EBS volumes. What will be the status of the alarms on the EBS volume?

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

Answer: B

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

NEW QUESTION 103
A user has launched an EC2 instance from an instance store backed AMI. The infrastructure team wants to create an AMI from the running instance. Which of the below mentioned credentials is not required while creating the AMI?

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

Answer: C

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

NEW QUESTION 104
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 107
An organization is measuring the latency of an application every minute and storing data inside a file in the JSON format. The organization wants to send all latency data to AWS CloudWatch. How can the organization achieve this?

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

Answer: C

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

NEW QUESTION 108
A user has launched an EBS backed instance with EC2-Classic. The user stops and starts the instance. Which of the below mentioned statements is not true with respect to the stop/start action?

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

Answer: C

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

NEW QUESTION 113
A user has launched multiple EC2 instances for the purpose of development and testing in the same region. The user wants to find the separate cost for the production and development instances. How can the user find the cost distribution?

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

Answer: D

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

NEW QUESTION 115
A user has created a VPC with CIDR 20.0.0.0/16 using VPC Wizard. The user has created a public CIDR (20.0.0.0/24. and a VPN only subnet CIDR (20.0.1.0/24. along with the hardware VPN access to connect to the user??s data centre. Which of the below mentioned components is not present when the VPC is setup with the wizard?

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

Answer: B

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

NEW QUESTION 116
A user has setup an Auto Scaling group. The group has failed to launch a single instance for more than 24 hours. What will happen to Auto Scaling in this condition?

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

Answer: B

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

NEW QUESTION 119
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 123
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 124
A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered. What time out period should the user specify for connection draining?

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

Answer: B

NEW QUESTION 127
A user is using the AWS EC2. The user wants to make so that when there is an issue in the EC2 server, such as instance status failed, it should start a new instance in the user??s private cloud. Which AWS service helps to achieve this automation?

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

D. AWS CloudWatch + AWS SNS

Answer: D

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

NEW QUESTION 128
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 131
A user has enabled session stickiness with ELB. The user does not want ELB to manage the cookie; instead he wants the application to manage the cookie. What will happen when the server instance, which is bound to a cookie, crashes?

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

Answer: B

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

NEW QUESTION 136
A user has created a VPC with public and private subnets using the VPC wizard. Which of the below mentioned statements is not true in this scenario?

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

Answer: A

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

NEW QUESTION 137
A user has configured ELB with a TCP listener at ELB as well as on the back-end instances. The user wants to enable a proxy protocol to capture the source and destination IP information in the header. Which of the below mentioned statements helps the user understand a proxy protocol with TCP configuration?

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

Answer: A

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

NEW QUESTION 140
A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. Which of the below mentioned statements is true with respect to this scenario?

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

Answer: D

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

NEW QUESTION 144
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 145
A user had aggregated the CloudWatch metric data on the AMI ID. The user observed some abnormal behaviour of the CPU utilization metric while viewing the last 2 weeks of data. The user wants to share that data with his manager. How can the user achieve this easily with the AWS console?

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

Answer: A

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

NEW QUESTION 148
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 151
A user has configured ELB with SSL using a security policy for secure negotiation between the client and load balancer. Which of the below mentioned SSL protocols is not supported by the security policy?

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

Answer: A

Explanation:
Elastic Load Balancing uses a Secure Socket Layer (SSL. negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. Elastic Load Balancing supports the following versions of the SSL protocol:
TLS 1.2
TLS 1.1
TLS 1.0
SSL 3.0
SSL 2.0

NEW QUESTION 153
A user has launched an EC2 Windows instance from an instance store backed AMI. The user has also set the Instance initiated shutdown behavior to stop. What will happen when the user shuts down the OS?

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

Answer: B

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

NEW QUESTION 157
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 159
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 162
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 167
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 170
A user has created a VPC with public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24 . The NAT instance ID is i-a12345. Which of the below mentioned entries are required in the main route table attached with the private subnet to allow instances to connect with the internet?

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

Answer: A

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

NEW QUESTION 171
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 173
An AWS account owner has setup multiple IAM users. One IAM user only has CloudWatch access. He has setup the alarm action which stops the EC2 instances when the CPU utilization is below the threshold limit. What will happen in this case?

A. It is not possible to stop the instance using the CloudWatch alarm
B. CloudWatch will stop the instance when the action is executed
C. The user cannot set an alarm on EC2 since he does not have the permission

D. The user can setup the action but it will not be executed if the user does not have EC2 rights

Answer: D

Explanation:
Amazon CloudWatch alarms watch a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup an action which stops 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. If the IAM user has read/write permissions for Amazon CloudWatch but not for Amazon EC2, he can still create an alarm. However, the stop or terminate actions will not be performed on the Amazon EC2 instance.

NEW QUESTION 175
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 by mistake. The user is trying to create another subnet of CIDR 20.0.0.1/24. How can the user create the second subnet?

A. There is no need to update the subnet as VPC automatically adjusts the CIDR of the first subnet based on the second subnet??s CIDR
B. The user can modify the first subnet CIDR from the console
C. It is not possible to create a second subnet as one subnet with the same CIDR as the VPC has been created
D. The user can modify the first subnet CIDR with AWS CLI

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 the 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. The user cannot modify the CIDR of a subnet once it is created. Thus, in this case if required, the user has to delete the subnet and create new subnets.

NEW QUESTION 176
A user has created a VPC with the public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The public subnet uses CIDR 20.0.1.0/24. The user is planning to host a web server in the public subnet (port 80. and a DB server in the private subnet (port 3306.. The user is configuring a security group for the public subnet (WebSecGrp. and the private subnet (DBSecGrp.. Which of the below mentioned entries is required in the web server security group (WebSecGrp.?

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

Answer: A

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

NEW QUESTION 177
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 182
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 187
A user has granted read/write permission of his S3 bucket using ACL. Which of the below mentioned options is a valid ID to grant permission to other AWS accounts (grantee. using ACL?

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

Answer: D

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

NEW QUESTION 188
A user has created an Auto Scaling group using CLI. The user wants to enable CloudWatch detailed monitoring for that group. How can the user configure this?

A. When the user sets an alarm on the Auto Scaling group, it automatically enables detail monitoring
B. By default detailed monitoring is enabled for Auto Scaling
C. Auto Scaling does not support detailed monitoring
D. Enable detail monitoring from the AWS console

Answer: B

Explanation:
CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. To enable detailed instance monitoring for a new Auto Scaling group, the user does not need to take any extra steps. When the user creates an Auto Scaling launch config as the first step for creating an Auto Scaling group, each launch configuration contains a flag named InstanceMonitoring.Enabled. The default value of this flag is true. Thus, the user does not need to set this flag if he wants detailed monitoring.

NEW QUESTION 189
A user has created a VPC with a public subnet. The user has terminated all the instances which are part of the subnet. Which of the below mentioned statements is true with respect to this scenario?

A. The user cannot delete the VPC since the subnet is not deleted
B. All network interface attached with the instances will be deleted
C. When the user launches a new instance it cannot use the same subnet
D. The subnet to which the instances were launched with will be deleted

Answer: B

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

NEW QUESTION 190
A user has created an application which will be hosted on EC2. The application makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK to connect with from the EC2 instance. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

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

Answer: A

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

NEW QUESTION 193
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 194
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 198
A user is trying to create a PIOPS EBS volume with 8 GB size and 200 IOPS. Will AWS create the volume?

A. Yes, since the ratio between EBS and IOPS is less than 30
B. No, since the PIOPS and EBS size ratio is less than 30
C. No, the EBS size is less than 10 GB
D. Yes, since PIOPS is higher than 100

Answer: C

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

NEW QUESTION 203
An organization has launched 5 instances: 2 for production and 3 for testing. The organization wants that one particular group of IAM users should only access the test instances and not the production ones. How can the organization set that as a part of the policy?

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

Answer: D

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The user can add conditions as a part of the IAM policies. The condition can be set on AWS Tags, Time, and Client IP as well as on various parameters. If the organization wants the user to access only specific instances he should define proper tags and add to the IAM policy condition. The sample policy is shown below.

```
"Statement": [
{
"Action": "ec2:*",
"Effect": "Allow",
"Resource": "*", "Condition": { "StringEquals": {
"ec2:ResourceTag/InstanceType": "Production"
```

}
}
}
]

NEW QUESTION 206
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 211
A user wants to upload a complete folder to AWS S3 using the S3 Management console. How can the user perform this activity?

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

Answer: D

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

NEW QUESTION 215
Which of the below mentioned AWS RDS logs cannot be viewed from the console for MySQL?

A. Error Log
B. Slow Query Log
C. Transaction Log
D. General Log

Answer: C

Explanation:
The user can view, download, and watch the database logs using the Amazon RDS console, the Command Line Interface (CLI., or the Amazon RDS API. For the MySQL RDS, the user can view the error log, slow querylog, and general logs. RDS does not support viewing the transaction logs.

NEW QUESTION 220
An organization has created 10 IAM users. The organization wants each of the IAM users to have access to a separate DyanmoDB table. All the users are added to the same group and the organization wants to setup a group level policy for this. How can the organization achieve this?

A. Define the group policy and add a condition which allows the access based on the IAM name
B. Create a DynamoDB table with the same name as the IAM user name and define the policy rule which grants access based on the DynamoDB ARN using a variable
C. Create a separate DynamoDB database for each user and configure a policy in the group based on the DB variable
D. It is not possible to have a group level policy which allows different IAM users to different DynamoDB Tables

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. AWS DynamoDB has only tables and the organization cannot makeseparate databases. The organization should create a table with the same name as the IAM user name and use the ARN of DynamoDB as part of the group policy. The sample policy is shown below:
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": ["dynamodb:*"],
"Resource": "arn:aws:dynamodb:region:account-number-without-hyphens:table/${aws:username}"
}
]
}

NEW QUESTION 224
A user has created a VPC with two subnets: one public and one private. The user is planning to run the patch update for the instances in the private subnet. How can the instances in the private subnet connect to the internet?

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

Answer: D

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

NEW QUESTION 229
A sys admin has enabled a log on ELB. Which of the below mentioned activities are not captured by the log?

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

Answer: B

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

NEW QUESTION 234
A user runs the command ??dd if=/dev/xvdf of=/dev/null bs=1M?? on an EBS volume created from a snapshot and attached to a Linux instance. Which of the below mentioned activities is the user performing with the step given above?

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

Answer: A

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

NEW QUESTION 235
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 240
Which method can be used to prevent an IP address block from accessing public objects in an S3 bucket?

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

Answer: A

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

NEW QUESTION 242
What would happen to an RDS (Relational Database Service) multi-Availability Zone deployment if the primary DB instance fails?

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

Answer: D

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

NEW QUESTION 247
How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

A. Simply create a new volume in the other AZ and specify the original volume as the source.
B. Detach the volume, then use the ec2-migrate-volume command to move it to another AZ.
C. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
D. Detach the volume and attach it to another EC2 instance in the other AZ.

Answer: C

Explanation:
Snapshots can be used to create multiple new EBS volumes, expand the size of a volume, or move volumes across Availability Zones.
See: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumes.html

NEW QUESTION 249
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 250
A company needs to monitor the read and write IOPs metrics for their AWS MySQL RDS instance and send real-time alerts to their operations team. Which AWS services can accomplish this? Choose 2 answers

A. Amazon Simple Email Service
B. Amazon CloudWatch
C. Amazon Simple Queue Service
D. Amazon Route 53
E. Amazon Simple Notification Service

Answer: BE

NEW QUESTION 254
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 255
Which features can be used to restrict access to data in S3? Choose 2 answers

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

Answer: AC

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

NEW QUESTION 259
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 263
A syslog Administrator is created additional Amazon EC2 instances and receive an Instancelimitexceeded error.
What is the cause of the issue and how can it be resolve?

A. The Administrator has requested too many instances at once and must request fewer instances in batches
B. The concurrent running instance limit has been reached and an EC2 limit increase request must be filed with AWS Support
C. AWS does not currently nave enough available capacity and a different instance type must be used
D. The Administrator must specify the maximum number of instances to be ?V created provisioning EC stances

Answer: B

Explanation:
EC2 Service Limits: AWS sets limits for these resources on a per-region basis.
If you are getting an InstanceLimitExceeded error when you try to launch an instance, you have reached your concurrent running instance limit. For new AWS accounts, the default limit is 20. If you need additional running instances, complete the form at Request to Increase Amazon EC2 Instance Limit.
By default, all AWS accounts have a limit of 20 running instances at any time per region. If you attempt to start another one, even if it already existed in the stopped state, you will receive this error message.
To resolve this issue, you can do any of the following: Stop one of your other running instances
Contact AWS support and request your running EC2 instances quota limit be raised.

NEW QUESTION 264
A Developers that an Amazon EC2 instance has failed. The developer reports that all the data was stored on the root volume is now gone.
What is the expiation for this issue?

A. The instance was using an Amazon EBS root volume
B. The instance was using Amazon S3 as the root volume
C. The instance was using an instance store root volume
D. The root volume with the data exists but needs to be re-attached

Answer: A

Explanation:
If your instance is ebs-backed, then you will not lose ebs root volume storage if you launched it with the "delete-on-termination" set to false. See the ec2-run-instances command for more information. Other ephemeral volumes will be lost when the instance is stopped/terminated.
If it is not ebs-backed, you will lose the root data when you terminate the instance (you cannot "stop" instance store instances).

NEW QUESTION 268
The Security tram is connect because the number of AWS identity and access Management (IAM) policies being in the environment is increasing. The tasked a SysOps Administrator to report on the number of IAM policies in use and use the total IAM policies.
Which AWS service should the Administrator use to check how current IAM policy compares to current limits?

A. MWS Trusted Advisor
B. Amazon Inspector
C. AWS Config
D. Organizations

Answer: C

Explanation:
AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This enables you to simplify compliance auditing, security analysis, change management, and operational troubleshooting.

NEW QUESTION 271
A company has created a separate AWS account for all development work to protect the production environment in this development account, developers have permission to manipulate IAM policies and roles. Corporate policies require that developers and blocked from accessing some services. What is the BEST way to grant the developers privileges in the development account while still complying with corporate policies?

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

Answer: A

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

NEW QUESTION 272
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 276
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 281
A company is migrating an application to AWS that requires access to a legacy system, which remain in the company's data centre. The application runs inside a VPC in the company's AWS account. The application must offer a consistent and low-latency response to its users
How can these requirements be met?

A. Create a software-based VPN connection between the Amazon VPC and the on-premises network
B. Create an AWS Direct Connect connection between AWS and the on-premises network and then use a private virtual interface

C. Create a hardware-based IPsec VPN connection between the VPC in AWS and the on-premises network
D. Create an overlay network by using third-party software and use that to connect the X/PC back to the on-premises network

Answer: B

Explanation:
Private Connectivity to your Amazon VPC. You can use AWS Direct Connect to establish a private virtual interface from your on-premise network directly to your Amazon VPC, providing you with a private, high bandwidth network connection between your network and your VPC.

NEW QUESTION 284
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 289
A SysOps Administrator has set up a new Application Load Balancer (ALB) in front of a pair of private web server in multiple Availability Zones. After deployment an updates CloudFormation template with many changes, user now goes to one web server only.
What is the NOST likely reason that the traffic is not being balanced between both servers?

A. The faulty is returning HTTP 200 has been removed.
B. Sticky session have been disabled in the ALB for the working sever.
C. The ALB using a custom ping path that is not found on the faulty server.
D. The web client are using HTTP/2, which is terminated at the ALB.

Answer: B

Explanation:
Until now, the behavior of load balancers has been to route each request independently to the Amazon EC2 instance with the least load. With the stickiness feature, you can configure the load balancer to bind user sessions to specific application instances. All requests coming from the user during the session will be sent to the same application instance. Elastic Load Balancing supports two mechanisms to provide session stickiness: load balancer-generated HTTP cookies, which allow browser-based session lifetimes, and application-generated HTTP cookies, which allow application- specific session lifetimes. You can learn more about this feature by visiting the ELB Developers Guide.

NEW QUESTION 290
A System Administrator is trying to identify why Put Object calls are not made from an Amazon EC2 instance to an Amazon bucket in the same region.
The instance is launched in a subnet with CIDR range 10.1.0.24 and 'Auto assign public IP set to yes. The instance profile tied to this instance has AmazonS3Access policy.
Security group rules for the instance:

The route table for the subnet in which this instance is launched

Destination	Target
10.0.0.0/16	local

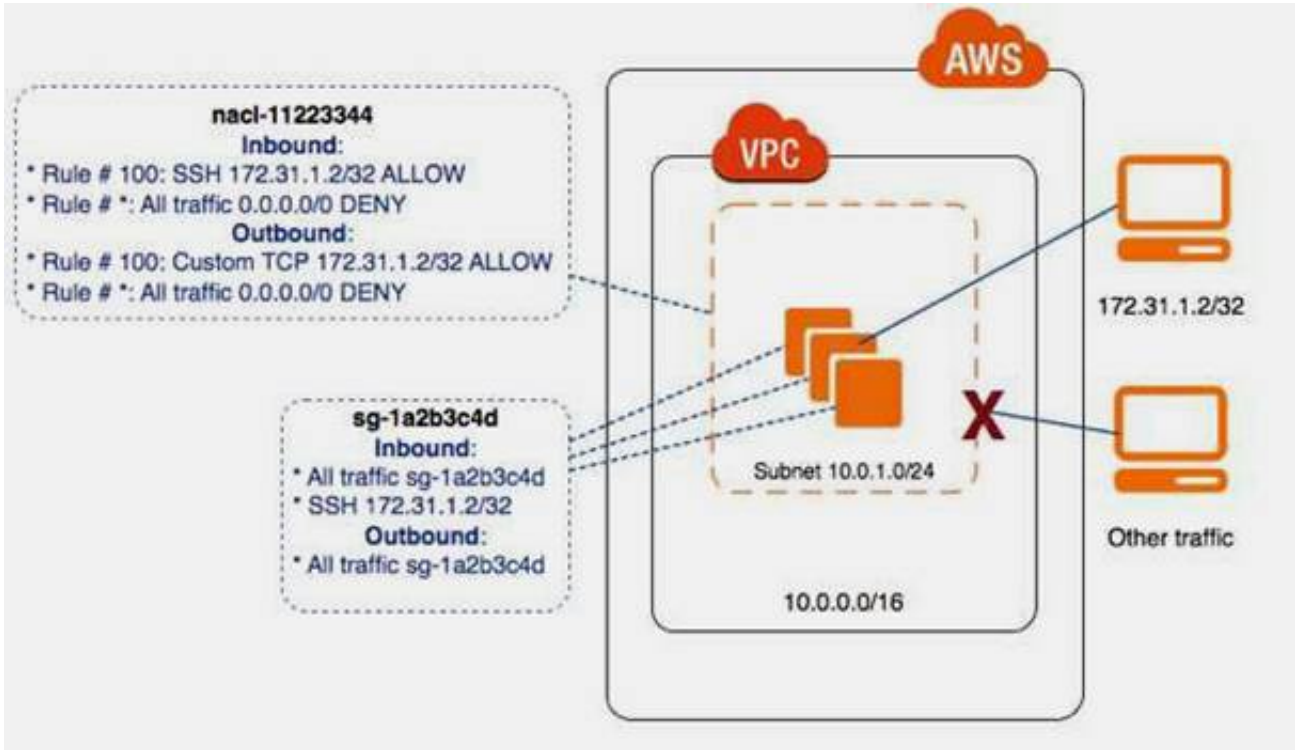Based on the information provided what is causing the lack of access to S3 from the instance?

A. The instance profile does not have explicit permissions to write objects to the S3 bucket.
B. The route table does not have a rule for all traffic to pass through a NAT gateway.
C. The route table does not have rule for all traffic to pass through an internet gateway

Answer: B

Explanation:
Controlling Access to Instances in a Subnet
In this example, instances in your subnet can communicate with each other, and are accessible from a trusted remote computer. The remote computer may be a computer in your local network or an instance in a different subnet or VPC that you use to connect to your instances to perform administrative tasks. Your security group rules and network ACL rules allow access from the IP address of your remote computer (172.31.1.2/32). All other traffic from the Internet or other networks is denied.

All instances use the same security group (sg-1a2b3c4d), with the following rules.
Protocol Protocol Port Source Comments

Type		Range		
All traffic	All	All	sg-1a2b3c4d	Enables instances associated with the same security group to communicate with each other.
TCP	SSH	22	172.31.1.2/32	Allows inbound SSH access from the remote computer. If the instance is a Windows computer, then this rule must use the RDP protocol for port 3389 instead.

Protocol Type	Protocol	Port Range	Destination	Comments
All traffic	All	All	sg-1a2b3c4d	Enables instances associated with the same security group to communicate with each other.

The subnet is associated with a network ACL that has the following rules.

Rule #	Type	Protocol	Port Range	Source	Allow/Deny	Comments
100	SSH	TCP	22	172.31.1.2/32	ALLOW	Allows inbound traffic from the remote computer. If the instance is a Windows computer, then this rule must use the RDP protocol for port 3389 instead.
*	All traffic All	All	All	0.0.0.0/0	DENY	Denies all other inbound traffic that does not match the previous rule.

Rule #	Type	Protocol	Port Range	Destination	Allow/Deny	Comments
100	Custom TCP	TCP	1024-65535	172.31.1.2/32	ALLOW	Allows outbound responses to the remote computer. Network ACLs are stateless, therefore this rule is required to allow response traffic for inbound requests.
*	All traffic All	All	All	0.0.0.0/0	DENY	Denies all other outbound traffic that does not match the previous rule.

This scenario gives you the flexibility to change the security groups or security group rules for your instances, and have the network ACL as the backup layer of defense. The network ACL rules apply to all instances in the subnet, so if you accidentally make your security group rules too permissive, the network ACL rules continue to permit access only from the single IP address. For example, the following rules are more permissive than the earlier rules ?X they allow inbound SSH access from any IP address.

Type	Protocol	Port Range	Source	Comments
All traffic	All	All	sg-1a2b3c4d	Enables instances associated with the same security group to communicate with each other.
SSH	TCP	22	0.0.0.0/0	Allows SSH access from any IP address.

Type	Protocol	Port Range	Destination	Comments
All traffic	All	All	0.0.0.0/0	Allows all outbound traffic.

However, only other instances within the subnet and your remote computer are able to access this instance. The network ACL rules still prevent all inbound traffic to the subnet except from your remote computer.

NEW QUESTION 292
An Application team is using Remote Desktop to connect to its application server and perform administrative tasks. After deployment a Windows service a existing subnets, the team discovers that it is unable to communicate with the new servers. A SysOps Administrative has obtained the VPC logs as shown in the table) related to the communication to help troubleshooting the problem.

How can this issue be resolved?

A. Check the route Tables to validate that the Remote Desktop and return traffic is allowed to and from the new servers.
B. Check the security groups to validate that Remote Desktop is allowed into the new servers.
C. Check the network access control lists to validate that the Remote Desktop and return traffic is allowed to and from the new servers.
D. Ensures that the RDP service and Windows firewall are open and listening on Port 3389 TCP.

Answer: D

NEW QUESTION 297
A company operate a secure website running an Amazon EC2 instance behind a Classic Load Balancer. An SSL certificate from AWS Certificate Manager is deployment on the load balancer. The company's Marketing team has determined that too many customer using older browser are experiencing issues with the website has asked a SysOps Administrator to fix this issue.
What course of action should the administrator take?

A. Update the SSL negotiation configuration of the load balancer by creating a custom security polic
B. Ensure the appropriate cipher has been enabled so that the web application can support the webbrowser.
C. Create a separate Classic Load Balancer and install custom SSL certificate with a different domain name on it that support the web browse
D. Ask customer with the affected browser to use this domain name instead of the one they are accustomed to using.
E. Create a new SSL certificate in Certificate Manager and install this certificate on each of the servers to accommodates the web browsers.
F. Remove the load balancer from the configuration and instead install a custom SSL certificate on each of the web servers.

Answer: A

Explanation:
Update the SSL Negotiation Configuration of Your Classic Load Balancer
Elastic Load Balancing provides security policies that have predefined SSL negotiation configurations to use to negotiate SSL connections between clients and your load balancer. If you are using the HTTPS/SSL protocol for your listener, you can use one of the predefined security policies, or use your own custom security policy.
For more information about the security policies, see SSL Negotiation Configurations for Classic Load Balancers. For information about the configurations of the security policies provided by Elastic Load Balancing, see Predefined SSL Security Policies.
If you create an HTTPS/SSL listener without associating a security policy, Elastic Load Balancing associates the default predefined security policy, ELBSecurityPolicy-2016-08, with your load balancer. If you have an existing load balancer with an SSL negotiation configuration that does not use the latest protocols and ciphers, we recommend that you update your load balancer to use ELBSecurityPolicy-2016-08. If you prefer, you can create a custom configuration. We strongly recommend that you test the new security policies before you upgrade your load balancer configuration.
The following examples show you how to update the SSL negotiation configuration for an HTTPS/SSL listener. Note that the change does not affect requests that were received by a load balancer node and are pending routing to a healthy instance, but the updated configuration will be used with new requests that are received.

NEW QUESTION 301
A company has mandated the use factor authentication (MFA) for all user, and requires users to make all API calls using CLI. However, uses are not prompted to enter MFA token, and able to return CLI commands without MF

A. In an enforce MFA, the company attached an IAM policy to all users that derives API calls that not been authenticated with MF
B. What additional step must be ensure that calls are authenticated using MFA?
C. Enable MFA on IAM roles, requires IAM to use role credentials to sign API calls.
D. Ask the IAM to log into the AWS Management Console with MFA before marking PI calls using the Cli.
E. Restricted the IAM users to use the console, as MFA not supported for CLI use.
F. Reporting users to use temporary credential from the get-session token command to sign API calls.

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/authenticate-mfa-cli/

NEW QUESTION 302
A SysOps Administrator has an AWS Lambda function that performs maintenance on versions AWS resources. This function must be run nightly.
Which is the MOST cost-effective solution?

A. Launch a single l2.nano Amazon EC2 instance and create a Linux corn job to invoke the Lambda function at the same every right.
B. Set up an Amazon CloudWatch metric alarm to invoke the Lambda function at the same time every night.
C. Schedule a CloudWatch event to invoke the Lambda function at the same time every night.
D. Implement a Chef recipe in Opsworks stack to invoke the Lambda function at the same time every night

Answer: C

Explanation:
Using AWS Lambda with Amazon CloudWatch Events
You can create a Lambda function and direct AWS Lambda to execute it on a regular schedule. You can specify a fixed rate (for example, execute a Lambda function every hour or 15 minutes), or you can specify a Cron expression. For more information on expressions schedules, see Schedule Expressions Using Rate or Cron.
This functionality is available when you create a Lambda function using the AWS Lambda console or the AWS CLI. To configure it using the AWS CLI, see Run an AWS Lambda Function on a Schedule Using the AWS CLI. The console provides CloudWatch Events as an event source. At the time of creating a Lambda function, you choose this event source and specify a time interval.
If you have made any manual changes to the permissions on your function, you may need to reapply
the scheduled event access to your function. You can do that by using the following CLI command.
$ aws lambda add-permission --function-name function_name\
--action 'lambda:InvokeFunction' --principal events.amazonaws.com \
--statement-id 'statement_id' \
--source-arn arn:aws:events:region:account-id:rule/rule_name

Each AWS account can have up to 100 unique event sources of the CloudWatch Events- Schedule source type. Each of these can be the event source for up to five Lambda functions. That is, you can have up to 500 Lambda functions that can be executing on a schedule in your AWS account.

The console also provides a blueprint (lambda-canary) that uses the CloudWatch Events - Schedule source type. Using this blueprint, you can create a sample Lambda function and test this feature. The example code that the blueprint provides checks for the presence of a specific webpage and specific text string on the webpage. If either the webpage or the text string is not found, the Lambda function throws an error.

NEW QUESTION 306
......

Relate Links

100% Pass Your SOA-C01 Exam with Exambible Prep Materials

https://www.exambible.com/SOA-C01-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/