



Amazon-Web-Services

Exam Questions SOA-C01

AWS Certified SysOps Administrator - Associate

NEW QUESTION 1

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

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

Answer: D

Explanation:

The question is about most ??scalable solution for communicating?? for SQS that is parallel processing of SQS messages.

See also:

?V <https://aws.amazon.com/articles/1464>

?V <http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/throughput.html>

NEW QUESTION 2

You have identified network throughput as a bottleneck on your m1.small EC2 instance when uploading data into Amazon S3 in the same region. How do you remedy this situation?

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 3

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

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

Answer: CD

NEW QUESTION 4

Which of the following are characteristics of Amazon VPC subnets? Choose 2 answers

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

Answer: AD

Explanation:

You can create a VPC that spans multiple Availability Zones. For more information, see [Creating a VPC](#). After creating a VPC, you can add one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span zones. Availability Zones are distinct locations that are engineered to be isolated from failures in other Availability Zones. By launching instances in separate Availability Zones, you can protect your applications from the failure of a single location. AWS assigns a unique ID to each subnet.

?V B is wrong: /28 is the smallest

?V C is wrong: private subnet should go via NAT (EIP only in public subnet)

?V E is wrong: subnet can only map to ONE AZ (not span multiple)

NEW QUESTION 5

You are creating an Auto Scaling group whose instances need to insert a custom metric into CloudWatch. Which method would be the best way to authenticate your CloudWatch PUT request?

- A. Create an IAM role with the Put MetricData permission and modify the Auto Scaling launch configuration to launch instances in that role
- B. Create an IAM user with the Put MetricData permission and modify the Auto Scaling launch configuration to inject the user's credentials into the instance User Data
- C. Modify the appropriate Cloud Watch metric policies to allow the Put MetricData permission to instances from the Auto Scaling group
- D. Create an IAM user with the Put MetricData permission and put the credentials in a private repository and have applications on the server pull the credentials as needed

Answer: A

NEW QUESTION 6

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 7

An application that you are managing has EC2 instances & Dynamo DB 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/GetSingleMetricAllDimensions.html>

NEW QUESTION 8

When assessing an organization s use of AWS API access credentials which of the following three credentials should be evaluated? Choose 3 answers

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

Answer: ACD

Explanation:

Reference:

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

NEW QUESTION 9

You have a Linux EC2 web server instance running inside a VPC The instance is In a public subnet and has an EIP associated with it so you can connect to It over the Internet via HTTP or SSH The instance was also fully accessible when you last logged in via SSH. and was also serving web requests on port 80.

Now you are not able to SSH into the host nor does it respond to web requests on port 80 that were working fine last time you checked You have double-checked that all networking configuration parameters (security groups route tables. IGW'EIP. NACLs etc) are properly configured {and you haven't made any changes to those anyway since you were last able to reach the Instance). You look at the EC2 console and notice that system status check shows "impaired."

Which should be your next step in troubleshooting and attempting to get the instance back to a healthy state so that you can log in again?

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

Answer: A

NEW QUESTION 10

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. Yes
- 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. Yes
- 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 10

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 12

You have a web-style application with a stateless but CPU and memory-intensive web tier running on a cc2 8xlarge EC2 instance inside of a VPC. The instance when under load is having problems returning requests within the SLA as defined by your business. The application maintains its state in a DynamoDB table, but the data tier is properly provisioned and responses are consistently fast.

How can you best resolve the issue of the application responses not meeting your SLA?

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

Answer: C

Explanation:

But it is possibly A as DynamoDB is automatically available across three facilities in an AWS Region. So moving in to a same AZ is not possible / necessary.

In this case the DB layer is not the issue, the EC2 8xlarge is the issue; so add another one with a ELB in-front of it.

See also: <https://aws.amazon.com/dynamodb/faqs/>

NEW QUESTION 13

You are designing a system that has a Bastion host. This component needs to be highly available without human intervention.

Which of the following approaches would you select?

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

Answer: C

NEW QUESTION 14

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 16

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

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

Answer: C

Explanation:

Reference:

<http://aws.amazon.com/cloudwatch/>

NEW QUESTION 19

You run a web application where web servers on EC2 Instances are In an Auto Scaling group Monitoring over the last 6 months shows that 6 web servers are necessary to handle the minimum load During the day up to 12 servers are needed Five to six days per year, the number of web servers required might go up to 15.

What would you recommend to minimize costs while being able to provide high availability?

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

Answer: A

Explanation:

The only plausible answer is A because all other answers include Spot Instances that can be removed without warning and that would not be highly available.

NEW QUESTION 20

You have set up Individual AWS accounts for each project. You have been asked to make sure your AWS Infrastructure costs do not exceed the budget set per project for each month.

Which of the following approaches can help ensure that you do not exceed the budget each month?

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

Answer: C

NEW QUESTION 23

You are using ElastiCache Memcached to store session state and cache database queries in your infrastructure. You notice in CloudWatch that Evictions and GetMisses are both very high.

What two actions could you take to rectify this? Choose 2 answers

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

Answer: AB

Explanation:

<http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheMetrics.WhichShouldIMonitor.html>

NEW QUESTION 27

You are running a database on an EC2 instance, with the data stored on Elastic Block Store (EBS) for persistence. At times throughout the day, you are seeing large variance in the response times of the database queries. Looking into the instance with the `iotop` command you see a lot of wait time on the disk volume that the database's data is stored on.

What two ways can you improve the performance of the database's storage while maintaining the current persistence of the data? Choose 2 answers

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

Answer: BC

NEW QUESTION 29

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 33

You are tasked with the migration of a highly trafficked Node JS application to AWS. In order to comply with organizational standards, Chef recipes must be used to configure the application servers that host this application and to support application lifecycle events.

Which deployment option meets these requirements while minimizing administrative burden?

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

Answer: A

Explanation:

OpsWorks has integrated support for Chef and lifecycle events.

See: <http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook.html>

NEW QUESTION 38

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 IAM users
- B. Create IAM users for privileged accounts
- C. Implement identity federation between your organization's Identity provider leveraging the IAM Security Token Service
- D. Enable the IAM 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 39

What are characteristics of Amazon S3? Choose 2 answers

- A. Objects are directly accessible via a URL
- B. S3 should be used to host a relational database
- C. S3 allows you to store objects or virtually unlimited size
- D. S3 allows you to store virtually unlimited amounts of data
- E. S3 offers Provisioned IOPS

Answer: AD

NEW QUESTION 40

You receive a frantic call from a new DBA who accidentally dropped a table containing all your customers.

Which Amazon RDS feature will allow you to reliably restore your database to within 5 minutes of when the mistake was made?

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

Answer: D

Explanation:

Reference:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.BackingUpAndRestoringAmazonRDSInstances.html>

NEW QUESTION 43

You use S3 to store critical data for your company. Several users within your group currently have full permissions to your S3 buckets. You need to come up with a solution that does not impact your users and also protect against the accidental deletion of objects.

Which two options will address this issue? Choose 2 answers

- A. Enable versioning on your S3 Buckets
- B. Configure your S3 Buckets with MFA delete
- C. Create a Bucket policy and only allow read only permissions to all users at the bucket level
- D. Enable object life cycle policies and configure the data older than 3 months to be archived in Glacier

Answer: AD

NEW QUESTION 48

An organization's security policy requires multiple copies of all critical data to be replicated across at least a primary and backup data center. The organization has decided to store some critical data on Amazon S3.

Which option should you implement to ensure this requirement is met?

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

Answer: D

Explanation:

It seems that this question wants to emphasize below (S3 Faq ?V <https://aws.amazon.com/s3/faqs/>) You specify a region when you create your Amazon S3 bucket. Within that region, your objects are redundantly stored on multiple devices across multiple facilities. Please refer to Regional Products and Services for details of Amazon S3 service availability by region.

NEW QUESTION 52

You are tasked with setting up a cluster of EC2 Instances for a NoSQL database. The database requires random read IO disk performance up to a 100,000 IOPS at 4KB block size per node.

Which of the following EC2 instances will perform the best for this workload?

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

Answer: C

Explanation:

Reference:

<http://aws.amazon.com/ec2/instance-types/>

NEW QUESTION 57

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 will 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 58

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

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

Answer: C

Explanation:

When you launch an instance into a VPC, a primary private IP address from the address range of the subnet is assigned to the default network interface (eth0) of the instance. If you don't specify a primary private IP address, we select an available IP address in the subnet range for you
<http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-ip-addressing.html>

NEW QUESTION 60

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

- A. The IP of the primary DB instance is switched to the standby DB 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 61

An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

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

Answer: C

Explanation:

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

NEW QUESTION 64

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

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

Answer: C

Explanation:

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

NEW QUESTION 68

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

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

Answer: C

Explanation:

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

NEW QUESTION 71

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 76

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 80

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 scale 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 81

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.html>

NEW QUESTION 86

A user is trying to save some cost on the AWS services. Which of the below mentioned options will not help him save cost?

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

Answer: B

Explanation:

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

NEW QUESTION 89

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 90

An organization is setting up programmatic billing access for their AWS account. Which of the below mentioned services is not required or enabled when the organization wants to use programmatic access?

- A. Programmatic access
- B. AWS bucket to hold the billing report
- C. AWS billing alerts
- D. Monthly Billing report

Answer: C

Explanation:

AWS provides an option to have programmatic access to billing. Programmatic Billing Access leverages the existing Amazon Simple Storage Service (Amazon S3) APIs. Thus, the user can build applications that reference his billing data from a CSV (comma-separated value) file stored in an Amazon S3 bucket. To enable programmatic access, the user has to first enable the monthly billing report. Then the user needs to provide an AWS bucket name where the billing CSV will be uploaded. The user should also enable the Programmatic access option.

NEW QUESTION 93

A user has configured the Auto Scaling group with the minimum capacity as 3 and the maximum capacity as 5. When the user configures the AS group, how many instances will Auto Scaling launch?

- A. 3
- B. 5
- C. 2

Answer: A

Explanation:

The default launch is going to configure 3 instances. It will only go to 5 under certain conditions. It won't launch initially 5 instances. This question doesn't give enough information to answer fully.

NEW QUESTION 95

An admin is planning to monitor the ELB. Which of the below mentioned services does not help the admin capture the monitoring information about the ELB activity?

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

Answer: B

Explanation:

The admin can capture information about Elastic Load Balancer using either:

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

NEW QUESTION 98

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

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

Answer: A

Explanation:

Auto Scaling has its own aggregated CPU Utilization metric.

NEW QUESTION 99

A user is trying to understand AWS SNS. To which of the below mentioned end points is SNS unable to send a notification?

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

Answer: D

Explanation:

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

NEW QUESTION 103

A user has configured an Auto Scaling group with ELB. The user has enabled detailed CloudWatch monitoring on Auto Scaling. Which of the below mentioned statements will help the user understand the functionality better?

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

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html CloudWatch monitors the following services. As soon as you begin using a service, it automatically sends metrics to CloudWatch for you.

CloudWatch offers either basic or detailed monitoring for supported AWS products. Basic monitoring means that a service sends data points to CloudWatch every five minutes. Detailed monitoring means that a service sends data points to CloudWatch every minute.

Note

If you are using a service that supports both basic and detailed data collection (for example, Amazon EC2 and Auto Scaling), and you want to access detailed statistics, you must enable detailed metric collection for that service.

Auto Scaling

Auto Scaling sends data to CloudWatch every 5 minutes by default. For an additional charge, you can enable detailed monitoring for Auto Scaling, which sends data to CloudWatch every minute. You can create alarms using Auto Scaling Dimensions and Metrics. For more information, see Monitor Your Auto Scaling Instances in the Auto Scaling User Guide.

NEW QUESTION 106

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

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

Answer: A

Explanation:

AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. While

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

NEW QUESTION 110

A system admin is managing buckets, objects and folders with AWS S3. Which of the below mentioned statements is true and should be taken in consideration by the sysadmin?

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

Answer: A

Explanation:

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

NEW QUESTION 113

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

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

Answer: D

Explanation:

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

NEW QUESTION 116

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

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

Answer: B

Explanation:

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

NEW QUESTION 120

A sys admin is maintaining an application on AWS. The application is installed on EC2 and user has configured ELB and Auto Scaling. Considering future load increase, the user is planning to launch new servers proactively so that they get registered with ELB. How can the user add these instances with Auto Scaling?

- A. Increase the desired capacity of the Auto Scaling group
- B. Increase the maximum limit of the Auto Scaling group
- C. Launch an instance manually and register it with ELB on the fly
- D. Decrease the minimum limit of the Auto Scaling group

Answer: A

Explanation:

A user can increase the desired capacity of the Auto Scaling group and Auto Scaling will launch a new instance as per the new capacity. The newly launched instances will be registered with ELB if Auto Scaling group is configured with ELB. If the user decreases the minimum size the instances will be removed from Auto Scaling. Increasing the maximum size will not add instances but only set the maximum instance cap.

NEW QUESTION 122

A root AWS account owner is trying to understand various options to set the permission to AWS S3. Which of the below mentioned options is not the right option to grant permission for S3?

- A. User Access Policy
- B. S3 Object Access Policy
- C. S3 Bucket Access Policy
- D. S3 ACL

Answer: B

Explanation:

Amazon S3 provides a set of operations to work with the Amazon S3 resources. Managing S3 resource access refers to granting others permissions to work with S3. There are three ways the root account owner can define access with S3:

S3 ACL: The user can use ACLs to grant basic read/write permissions to other AWS accounts.

S3 Bucket Policy: The policy is used to grant other AWS accounts or IAM users permissions for the bucket and the objects in it.

User Access Policy: Define an IAM user and assign him the IAM policy which grants him access to S3.

NEW QUESTION 127

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 128

A user has enabled the Multi AZ feature with the MS SQL RDS database server. Which of the below mentioned statements will help the user understand the Multi AZ feature better?

- A. In a Multi AZ, AWS runs two DBs in parallel and copies the data asynchronously to the replica copy
- B. In a Multi AZ, AWS runs two DBs in parallel and copies the data synchronously to the replica copy
- C. In a Multi AZ, AWS runs just one DB but copies the data synchronously to the standby replica
- D. AWS MS SQL does not support the Multi AZ feature

Answer: C

Explanation:

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy, eliminate I/O freezes, and minimize latency spikes during system backups. Running a DB instance with high availability can enhance availability during planned system maintenance, and help protect your databases against DB instance failure and Availability Zone disruption. Note that the high-availability feature is not a scaling solution for read-only scenarios; you cannot use a standby replica to serve read traffic. To service read-only traffic, you should use a read replica.

NEW QUESTION 130

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

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

Answer: D

Explanation:

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

NEW QUESTION 131

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

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

Answer: B

Explanation:

When launching an instance with EC2, AWS recommends not to select the availability zone (AZ). AWS specifies that the default Availability Zone should be accepted. This is because it enables AWS to select the best Availability Zone based on the system health and available capacity. If the user launches additional

instances, only then an Availability Zone should be specified. This is to specify the same or different AZ from the running instances.

NEW QUESTION 134

A user has created a VPC with CIDR 20.0.0.0/16 with only a private subnet and VPN connection using the VPC wizard. The user wants to connect to the instance in a private subnet over SSH. How should the user define the security rule for SSH?

- A. Allow Inbound traffic on port 22 from the user's network
- B. The user has to create an instance in EC2 Classic with an elastic IP and configure the security group of a private subnet to allow SSH from that elastic IP
- C. The user can connect to a instance in a private subnet using the NAT instance
- D. Allow Inbound traffic on port 80 and 22 to allow the user to connect to a private subnet over the Internet

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 center, the user can setup a case with a VPN only subnet (private. which uses VPN access to connect with his data center. When the user has configured this setup with Wizard, all network connections to the instances in the subnet will come from his data center. The user has to configure the security group of the private subnet which allows the inbound traffic on SSH (port 22. from the data center's network range.

NEW QUESTION 136

A user has enabled detailed CloudWatch metric monitoring on an Auto Scaling group. Which of the below mentioned metrics will help the user identify the total number of instances in an Auto Scaling group cluding pending, terminating and running instances?

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

Answer: A

Explanation:

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

NEW QUESTION 141

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 143

A user is checking the CloudWatch metrics from the AWS console. The user notices that the CloudWatch data is coming in UTC. The user wants to convert the data to a local time zone. How can the user perform this?

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

Answer: B

Explanation:

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

NEW QUESTION 147

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 151

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

- A. <http://sqs.us-east-1.amazonaws.com/123456789012/myqueue>
- B. <http://sqs.amazonaws.com/123456789012/myqueue>
- C. <http://sq>
- D. 123456789012.us-east-1.amazonaws.com/myqueue
- E. [http:// 123456789012.sq](http://123456789012.sq)
- F. us-east-1.amazonaws.com/myqueue

Answer: A

Explanation:

When creating a new queue in SQS, the user must provide a queue name that is unique within the scope of all queues of user??s account. If the user creates queues using both the latest WSDL and a previous version, he will have a single namespace for all his queues. Amazon SQS assigns each queue created by user an identifier called a queue URL, which includes the queue name and other components that Amazon SQS determines. Whenever the user wants to perform an action on a queue, he must provide its queue URL. The queue URL for the account id 123456789012 & queue name ??myqueue?? in US-East-1 region will be <http://sqs.us-east-1.amazonaws.com/123456789012/myqueue>.

NEW QUESTION 155

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

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

Answer: A

Explanation:

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

NEW QUESTION 158

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

- A. The user should not use RI; instead only go with the on-demand pricing
- B. The user should use the AWS high utilized RI
- C. The user should use the AWS medium utilized RI
- D. The user should use the AWS low utilized RI

Answer: A

Explanation:

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

NEW QUESTION 160

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 165

A user has launched an EBS backed EC2 instance. The user has rebooted the instance. Which of the below mentioned statements is not true with respect to the

reboot action?

- A. The private and public address remains the same
- B. The Elastic IP remains associated with the instance
- C. The volume is preserved
- D. The instance runs on a new host computer

Answer: D

Explanation:

A user can reboot an EC2 instance using the AWS console, the Amazon EC2 CLI or the Amazon EC2 API. Rebooting an instance is equivalent to rebooting an operating system. However, it is recommended that the user use the Amazon EC2 to reboot the instance instead of running the operating system reboot command from the instance. The instance remains on the same host computer and maintains its public DNS name, private IP address, and any data on its instance store volumes. It typically takes a few minutes for the reboot to complete, but the time it takes to reboot depends on the instance configuration.

NEW QUESTION 169

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

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

Answer: C

Explanation:

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

NEW QUESTION 172

A user has created a public subnet with VPC and launched an EC2 instance within it. The user is trying to delete the subnet. What will happen in this scenario?

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

Answer: B

Explanation:

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

NEW QUESTION 175

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 177

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 179

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) to organize usage costs across multiple services.

NEW QUESTION 182

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 186

A user is planning to set up the Multi AZ feature of RDS. Which of the below mentioned conditions won't take advantage of the Multi AZ feature?

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

Answer: C

Explanation:

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

NEW QUESTION 189

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

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

Answer: B

Explanation:

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

NEW QUESTION 191

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

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

Answer: C

Explanation:

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

NEW QUESTION 196

A sysadmin has created the below mentioned policy on an S3 bucket named cloudacademy. What does this policy define?

```
"Statement": [{
  "Sid": "Stmnt1388811069831",
  "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 about the 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 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": "Stmnt1388811069831",
  "Effect": "Allow", "Principal": { "AWS": "*" },
  "Action": [ "s3:GetObjectAcl", "s3:ListBucket"], "Resource": [ "arn:aws:s3:::cloudacademy"]
}]
```

NEW QUESTION 200

An organization (account ID 123412341234) has configured the IAM policy to allow the user to modify his credentials. What will the below mentioned statement allow the user to perform?

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow", "Action": [ "iam:AddUserToGroup",
    "iam:RemoveUserFromGroup", "iam:GetGroup"
    ],
    "Resource": "arn:aws:iam:: 123412341234:group/TestingGroup"
  }]
}
```

- A. The IAM policy will throw an error due to an invalid resource name
- B. The IAM policy will allow the user to subscribe to any IAM group
- C. Allow the IAM user to update the membership of the group called TestingGroup
- D. Allow the IAM user to delete the TestingGroup

Answer: C

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (account ID 123412341234) wants their users to manage their subscription to the groups, they should create a relevant policy for that. The below mentioned policy allows the respective IAM user to update the membership of the group called MarketingGroup.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow", "Action": [ "iam:AddUserToGroup",
    "iam:RemoveUserFromGroup", "iam:GetGroup"
    ],
    "Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup "
  }]
}
```

NEW QUESTION 205

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

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

Answer: A

Explanation:

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

NEW QUESTION 206

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 207

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 212

An organization has setup multiple IAM users. The organization wants that each IAM user accesses the IAM console only within the organization and not from outside. How can it achieve this?

- A. Create an IAM policy with the security group and use that security group for AWS console login
- B. Create an IAM policy with a condition which denies access when the IP address range is not from the organization
- C. Configure the EC2 instance security group which allows traffic only from the organization's IP range
- D. Create an IAM policy with VPC and allow a secure gateway between the organization and AWS Console

Answer: B

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 many other parameters. If the organization wants the user to access only from a specific IP range, they should set an IAM policy condition which denies access when the IP is not in a certain range. E.g. The sample policy given below denies all traffic when the IP is not in a certain range.

```
"Statement": [{  
  "Effect": "Deny",  
  "Action": "*",  
  "Resource": "*", "Condition": { "NotIpAddress": {  
    "aws:SourceIp": ["10.10.10.0/24", "20.20.30.0/24"]  
  }  
}  
}]
```

NEW QUESTION 217

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 222

A user is observing the EC2 CPU utilization metric on CloudWatch. The user has observed some interesting patterns while filtering over the 1 week period for a

particular hour. The user wants to zoom that data point to a more granular period. How can the user do that easily with CloudWatch?

- A. The user can zoom a particular period by selecting that period with the mouse and then releasing the mouse
- B. The user can zoom a particular period by double clicking on that period with the mouse
- C. The user can zoom a particular period by specifying the aggregation data for that period
- D. The user can zoom a particular period by specifying the period in the Time Range

Answer: A

NEW QUESTION 223

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

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

Answer: B

Explanation:

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

NEW QUESTION 226

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 229

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 231

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

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

Answer: A

Explanation:

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

NEW QUESTION 236

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

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

Answer: D

Explanation:

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

NEW QUESTION 240

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 known 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 241

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 244

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

- A. The ratio between IOPS and the EBS volume is higher than 30
- B. The maximum IOPS supported by EBS is 3000

- C. The ratio between IOPS and the EBS volume is lower than 50
- D. PIOPS is supported for EBS higher than 500 GB size

Answer: A

Explanation:

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

NEW QUESTION 245

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

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

Answer: B

NEW QUESTION 246

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 251

A user has enabled versioning on an S3 bucket. The user is using server side encryption for data at Rest. If the user is supplying his own keys for encryption (SSE-C), which of the below mentioned statements is true?

- A. The user should use the same encryption key for all versions of the same object
- B. It is possible to have different encryption keys for different versions of the same object
- C. AWS S3 does not allow the user to upload his own keys for server side encryption
- D. The SSE-C does not work when versioning is enabled

Answer: B

Explanation:

AWS S3 supports client side or server side encryption to encrypt all data at rest. The server side encryption can either have the S3 supplied AES-256 encryption key or the user can send the key along with each API call to supply his own encryption key (SSE-C). If the bucket is versioning-enabled, each object version uploaded by the user using the SSE-C feature can have its own encryption key. The user is responsible for tracking which encryption key was used for which object's version

NEW QUESTION 253

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 254

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 258

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 262

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

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

Answer: C

Explanation:

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

NEW QUESTION 266

A user is configuring the Multi AZ feature of an RDS DB. The user came to know that this RDS DB does not use the AWS technology, but uses server mirroring to achieve H

- A. Which DB is the user using right now?
- B. My SQL
- C. Oracle
- D. MS SQL
- E. PostgreSQL

Answer: C

Explanation:

Amazon RDS provides high availability and failover support for DB instances using Multi AZ deployments. In a Multi AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. Multi AZ deployments for Oracle, PostgreSQL, and MySQL DB instances use Amazon technology, while SQL Server (MS SQL) DB instances use SQL Server Mirroring.

NEW QUESTION 269

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 274

A user is trying to understand the CloudWatch metrics for the AWS services. It is required that the user should first understand the namespace for the AWS services. Which of the below mentioned is not a valid namespace for the AWS services?

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

Answer: B

Explanation:

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

NEW QUESTION 277

A user has launched an EC2 instance from an instance store backed AMI. If the user restarts the instance, what will happen to the ephemeral storage data?

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

Answer: D

Explanation:

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

NEW QUESTION 278

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

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

Answer: B

Explanation:

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

NEW QUESTION 283

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

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

Answer: D

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. For ELB, CloudWatch provides various metrics including error code by ELB as well as by back-end servers (instances.. It gives data for the count of the number of HTTP response codes generated by the back-end instances. This metric does not include any response codes generated by the load balancer. These metrics are:

The 2XX class status codes represents successful actions

The 3XX class status code indicates that the user agent requires action The 4XX class status code represents client errors

The 5XX class status code represents back-end server errors

NEW QUESTION 287

A user has launched an EC2 instance store backed instance in the US-East-1a zone. The user created AMI #1 and copied it to the Europe region. After that, the user made a few updates to the application running in the US-East-1a zone. The user makes an AMI#2 after the changes. If the user launches a new instance in Europe from the AMI #1 copy, which of the below mentioned statements is true?

- A. The new instance will have the changes made after the AMI copy as AWS just copies the reference of the original AMI during the copyin
- B. Thus, the copied AMI will have all the updated data
- C. The new instance will have the changes made after the AMI copy since AWS keeps updating the AMI
- D. It is not possible to copy the instance store backed AMI from one region to another
- E. The new instance in the EU region will not have the changes made after the AMI copy

Answer: D

Explanation:

Within EC2, when the user copies an AMI, the new AMI is fully independent of the source AMI; there is no link to the original (source) AMI. The user can modify the source AMI without affecting the new AMI and vice versa. Therefore, in this case even if the source AMI is modified, the copied AMI of the EU region will not have the changes. Thus, after copy the user needs to copy the new source AMI to the destination region to get those changes.

NEW QUESTION 292

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 296

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 sign 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 301

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

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

Answer: A

Explanation:

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

```
chmod 0400 /path/to/private.key
```

NEW QUESTION 305

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

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

D. The volume size is too large

Answer: D

Explanation:

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

NEW QUESTION 307

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

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

Answer: D

Explanation:

Amazon RDS performs maintenance on the DB instance during a user-definable maintenance window. The system may be offline or experience lower performance during that window. The only maintenance events that may require RDS to make the DB instance offline are:

Scaling compute operations

Software patching. Required software patching is automatically scheduled only for patches that are security and durability related. Such patching occurs infrequently (typically once every few months, and seldom requires more than a fraction of the maintenance window.

NEW QUESTION 309

A user is trying to create an EBS volume with the highest PIOPS supported by EBS. What is the minimum size of EBS required to have the maximum IOPS?

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

Answer: C

Explanation:

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

NEW QUESTION 314

An organization is trying to create various IAM users. Which of the below mentioned options is not a valid IAM username?

- A. John.cloud
- B. john@cloud
- C. John=cloud
- D. john#cloud

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. Whenever the organization is creating an IAM user, there should be a unique ID for each user. The names of users, groups, roles, instance profiles must be alphanumeric, including the following common characters: plus (+), equal (=), comma (,), period (.), at (@), and dash (-).

NEW QUESTION 318

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

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

Answer: A

Explanation:

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

NEW QUESTION 322

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

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

Answer: B

Explanation:

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

NEW QUESTION 325

A user has two EC2 instances running in two separate regions. The user is running an internal memory management tool, which captures the data and sends it to CloudWatch in US East, using a CLI with the same namespace and metric. Which of the below mentioned options is true with respect to the above statement?

- A. The setup will not work as CloudWatch cannot receive data across regions
- B. CloudWatch will receive and aggregate the data based on the namespace and metric
- C. CloudWatch will give an error since the data will conflict due to two sources
- D. CloudWatch will take the data of the server, which sends the data first

Answer: B

Explanation:

Amazon CloudWatch does not differentiate the source of a metric when receiving custom data. If the user is publishing a metric with the same namespace and dimensions from different sources, CloudWatch will treat them as a single metric. If the data is coming with the same timezone within a minute, CloudWatch will aggregate the data. It treats these as a single metric, allowing the user to get the statistics, such as minimum, maximum, average, and the sum of all across all servers.

NEW QUESTION 326

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

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

Answer: D

Explanation:

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

NEW QUESTION 331

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

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

Answer: A

Explanation:

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

NEW QUESTION 332

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

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

Answer: A

Explanation:

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

manually.

NEW QUESTION 334

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 336

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

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

Answer: B

Explanation:

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

NEW QUESTION 341

A user has created a VPC with public and private subnets using the VPC Wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24. Which of the below mentioned entries are required in the main route table to allow the instances in VPC to communicate with each other?

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

Answer: D

NEW QUESTION 344

A sysadmin has created the below mentioned policy on an S3 bucket named cloudacademy. The bucket has both AWS.jpg and index.html objects. What does this policy define?

```
"Statement": [{
  "Sid": "Stmnt1388811069831",
  "Effect": "Allow", "Principal": { "AWS": "*" },
  "Action": [ "s3:GetObjectAcl", "s3:ListBucket", "s3:GetObject"], "Resource": [ "arn:aws:s3:::cloudacademy/* .jpg"]
}]
```

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

Answer: B

NEW QUESTION 345

A user has launched an EC2 instance and deployed a production application in it. The user wants to prohibit any mistakes from the production team to avoid accidental termination. How can the user achieve this?

- A. The user can the set DisableApiTermination attribute to avoid accidental termination
- B. It is not possible to avoid accidental termination
- C. The user can set the Deletion termination flag to avoid accidental termination
- D. The user can set the InstanceInitiatedShutdownBehavior flag to avoid accidental termination

Answer: A

Explanation:

It is always possible that someone can terminate an EC2 instance using the Amazon EC2 console, command line interface or API by mistake. If the admin wants to prevent the instance from being accidentally terminated, he can enable termination protection for that instance. The DisableApiTermination attribute controls whether the instance can be terminated using the console, CLI or API. By default, termination protection is disabled for an EC2 instance. When it is set it will not allow the user to terminate the instance from CLI, API or the console.

NEW QUESTION 347

A user has deployed an application on an EBS backed EC2 instance. For a better performance of application, it requires dedicated EC2 to EBS traffic. How can the user achieve this?

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

Answer: D

Explanation:

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

NEW QUESTION 349

A user has launched a Windows based EC2 instance. However, the instance has some issues and the user wants to check the log. When the user checks the Instance console output from the AWS console, what will it display?

- A. All the event logs since instance boot
- B. The last 10 system event log error
- C. The Windows instance does not support the console output
- D. The last three system events?? log errors

Answer: D

Explanation:

The AWS EC2 console provides a useful tool called Console output for problem diagnosis. It is useful to find out any kernel issues, termination reasons or service configuration issues. For a Windows instance it lists the last three system event log errors. For Linux it displays the exact console output.

NEW QUESTION 352

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

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

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

Answer: C

NEW QUESTION 353

Which services allow the customer to retain run administrative privileges or the underlying EC2 instances? Choose 2 answers

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

Answer: AB

NEW QUESTION 357

Amazon EBS snapshots have which of the following two characteristics? Choose 2 answers

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

Answer: AB

NEW QUESTION 361

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

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

Answer: ABE

NEW QUESTION 365

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 367

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SOA-C01 Practice Exam Features:

- * SOA-C01 Questions and Answers Updated Frequently
- * SOA-C01 Practice Questions Verified by Expert Senior Certified Staff
- * SOA-C01 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * SOA-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SOA-C01 Practice Test Here](#)