



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 an Auto Scaling group associated with an Elastic Load Balancer (ELB). You have noticed that instances launched via the Auto Scaling group are being marked unhealthy due to an ELB health check, but these unhealthy instances are not being terminated.

What do you need to do to ensure trial instances marked unhealthy by the ELB will be terminated and replaced?

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

Answer: B

Explanation:

Reference:

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-add-elb-healthcheck.html>

Add an Elastic Load Balancing Health Check to your Auto Scaling Group

By default, an Auto Scaling group periodically reviews the results of EC2 instance status to determine the health state of each instance. However, if you have associated your Auto Scaling group with an Elastic Load Balancing load balancer, you can choose to use the Elastic Load Balancing health check. In this case, Auto Scaling determines the health status of your instances by checking the results of both the EC2 instance status check and the Elastic Load Balancing instance health check.

For information about EC2 instance status checks, see Monitor Instances With Status Checks in the Amazon EC2 User Guide for Linux Instances. For information about Elastic Load Balancing health checks, see Health Check in the Elastic Load Balancing Developer Guide.

This topic shows you how to add an Elastic Load Balancing health check to your Auto Scaling group, assuming that you have created a load balancer and have registered the load balancer with your Auto Scaling group. If you have not registered the load balancer with your Auto Scaling group, see Set Up a Scaled and Load-Balanced Application.

Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action DescribeInstanceStatus return any state other than running, the system status shows impaired, or the calls to Elastic Load Balancing action DescribeInstanceHealth returns OutOfService in the instance state field.

If there are multiple load balancers associated with your Auto Scaling group, Auto Scaling checks the health state of your EC2 instances by making health check calls to each load balancer. For each call, if the Elastic Load Balancing action returns any state other than InService, the instance is marked as unhealthy. After Auto Scaling marks an instance as unhealthy, it remains in that state, even if subsequent calls from other load balancers return an InService state for the same instance.

NEW QUESTION 3

An organization has configured a VPC with an Internet Gateway (IGW). pairs of public and private subnets (each with one subnet per Availability Zone), and an Elastic Load Balancer (ELB) configured to use the public subnets. The application??s web tier leverages the ELB. Auto Scaling and a multi-AZ RDS database instance The organization would like to eliminate any potential single points of failure in this design.

What step should you take to achieve this organization's objective?

- A. Nothing, there are no single points of failure in this architecture.
- B. Create and attach a second IGW to provide redundant internet connectivity.
- C. Create and configure a second Elastic Load Balancer to provide a redundant load balancer.
- D. Create a second multi-AZ RDS instance in another Availability Zone and configure replication to provide a redundant database.

Answer: A

NEW QUESTION 4

When an EC2 instance that is backed by an S3-based AMI Is terminated, what happens to the data on me root volume?

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

Answer: C

Explanation:

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ComponentsAMIs.html>

NEW QUESTION 5

You have a web application leveraging an Elastic Load Balancer (ELB) in front of the web servers deployed using an Auto Scaling Group. Your database is running on Relational Database Service (RDS). The application serves out technical articles and responses to them in general there are more views of an article than there are responses to the article. On occasion, an article on the site becomes extremely popular resulting in significant traffic increases that causes the site to go down.

What could you do to help alleviate the pressure on the infrastructure while maintaining availability during these events?

Choose 3 answers

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

Answer: ABC

Explanation:

The questions mention RDS so an answer that includes that as part of the solution makes sense. Also, Route53 does nothing to alleviate pressure on the infrastructure, it's for failover. E is counterproductive. It talks about failing over to an error page on S3.

NEW QUESTION 6

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 7

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

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

Answer: BD

Explanation:

This is an odd question. First of all, option A cannot be right because ELB does not failover. Cannot be C because Route 53 does work with hard-coded IP. Only B & D cannot be ruled out so best answer.

NEW QUESTION 8

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 9

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 10

You have been asked to propose a multi-region deployment of a web-facing application where a controlled portion of your traffic is being processed by an alternate region.

Which configuration would achieve that goal?

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

Answer: A

Explanation:

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

NEW QUESTION 10

You have 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 14

When creation of an EBS snapshot is initiated but not completed the EBS volume?

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

Answer: C

Explanation:

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

NEW QUESTION 15

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 19

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 m2.4xlarge Instance Instead

Answer: BC

NEW QUESTION 22

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 24

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

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

Answer: B

Explanation:

Reference:
<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/UserScenariosForVPC.html>

NEW QUESTION 26

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 27

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

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

Answer: A

NEW QUESTION 28

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 31

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 34

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 36

A customer has a web application that uses cookie Based sessions to track logged in users It Is deployed on AWS using ELB and Auto Scaling The customer observes that when load increases. Auto Scaling launches new Instances but the load on the existing Instances does not decrease, causing all existing users to have a sluggish experience.

Which two answer choices independently describe a behavior that could be the cause of the sluggish user experience? Choose 2 answers

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

Answer: BD

NEW QUESTION 37

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

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

Answer: D

Explanation:

<https://aws.amazon.com/rds/faqs/>

NEW QUESTION 42

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 45

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

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

Answer: D

Explanation:

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

NEW QUESTION 48

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

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

Answer: D

Explanation:

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

NEW QUESTION 52

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 54

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 57

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 61

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

- A. Add a new instance manually by 8 AM Thursday and terminate the same by 6 PM Friday
- B. Schedule Auto Scaling to scale up by 8 AM Thursday and scale down after 6 PM on Friday
- C. Schedule a policy which may scale up every day at 8 AM and scales down by 6 PM
- D. Configure a batch process to add an 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 63

A user has setup a billing alarm using CloudWatch for \$200. The usage of AWS exceeded \$200 after some days. The user wants to increase the limit from \$200 to \$400. What should the user do?

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

Answer: C

Explanation:

AWS CloudWatch supports enabling the billing alarm on the total AWS charges. The estimated charges are calculated and sent several times daily to CloudWatch in the form of metric data. This data will be stored for 14 days. This data also includes the estimated charges for every service in AWS used by the user, as well as the estimated overall AWS charges. If the user wants to increase the limit, the user can modify the alarm and specify a new threshold.

NEW QUESTION 65

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 66

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 67

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 72

You are building an online store on AWS that uses SQS to process your customer orders. Your backend system needs those messages in the same sequence the customer orders have been put in. How can you achieve that?

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

Answer: B

Explanation:

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

NEW QUESTION 76

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 78

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 83

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

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

Answer: C

Explanation:

The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that in-flight requests continue to be served. The user can specify a maximum time (3600 seconds) for the load

balancer to keep the connections alive before reporting the instance as deregistered. If the user does not specify the maximum timeout period, by default, the load balancer will close the connections to the deregistering instance after 300 seconds.

NEW QUESTION 84

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 87

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

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

Answer: C

Explanation:

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

NEW QUESTION 92

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

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

Answer: A

Explanation:

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

NEW QUESTION 94

A user is planning to setup notifications on the RDS DB for a snapshot. Which of the below mentioned event categories is not supported by RDS for this snapshot source type?

- A. Backup
- B. Creation
- C. Deletion
- D. Restoration

Answer: A

Explanation:

Amazon RDS uses the Amazon Simple Notification Service to provide a notification when an Amazon RDS event occurs. Event categories for a snapshot source type include: Creation, Deletion, and Restoration. The Backup is a part of DB instance source type.

NEW QUESTION 96

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 100

An organization has configured the custom metric upload with CloudWatch. The organization has given permission to its employees to upload data using CLI as well SDK. How can the user track the calls made to CloudWatch?

- A. The user can enable logging with CloudWatch which logs all the activities
- B. Use CloudTrail to monitor the API calls
- C. Create an IAM user and allow each user to log the data using the S3 bucket
- D. Enable detailed monitoring with CloudWatch

Answer: B

Explanation:

AWS CloudTrail is a web service which will allow the user to monitor the calls made to the Amazon CloudWatch API for the organization's account, including calls made by the AWS Management Console, Command Line Interface (CLI), and other services. When CloudTrail logging is turned on, CloudWatch will write log files into the Amazon S3 bucket, which is specified during the CloudTrail configuration.

NEW QUESTION 103

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

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

Answer: B

Explanation:

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

NEW QUESTION 106

A user has stored data on an encrypted EBS volume. The user wants to share the data with his friend's AWS account. How can user achieve this?

- A. Create an AMI from the volume and share the AMI
- B. Copy the data to an unencrypted volume and then share
- C. Take a snapshot and share the snapshot with a friend
- D. If both the accounts are using the same encryption key then the user can share the volume directly

Answer: B

Explanation:

AWS EBS supports encryption of the volume. It also supports creating volumes from existing snapshots provided the snapshots are created from encrypted volumes. If the user is having data on an encrypted volume and is trying to share it with others, he has to copy the data from the encrypted volume to a new unencrypted volume. Only then can the user share it as an encrypted volume data. Otherwise the snapshot cannot be shared.

NEW QUESTION 107

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

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

Answer: B

Explanation:

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

NEW QUESTION 111

A user has created an ELB with the availability zone US-East-1

- A. The user wants to add more zones to ELB to achieve High Availability
- B. How can the user add more zones to the existing ELB?
- C. It is not possible to add more zones to the existing ELB

- D. The only option is to launch instances in different zones and add to ELB
- E. The user should stop the ELB and add zones and instances as required
- F. The user can add zones on the fly from the AWS console

Answer: D

Explanation:

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

NEW QUESTION 116

A user has configured CloudWatch monitoring on an EBS backed EC2 instance. If the user has not attached any additional device, which of the below mentioned metrics will always show a 0 value?

- A. DiskReadBytes
- B. NetworkIn
- C. NetworkOut
- D. CPUUtilization

Answer: A

Explanation:

CloudWatch is used to monitor AWS as the well custom services. For EC2 when the user is monitoring the EC2 instances, it will capture the 7 Instance level and 3 system check parameters for the EC2 instance. Since this is an EBS backed instance, it will not have ephemeral storage attached to it. Out of the 7 EC2 metrics, the 4 metrics DiskReadOps, DiskWriteOps, DiskReadBytes and DiskWriteBytes are disk related data and available only when there is ephemeral storage attached to an instance. For an EBS backed instance without any additional device, this data will be 0.

NEW QUESTION 117

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 119

A root account owner has created an S3 bucket testmycloud. The account owner wants to allow everyone to upload the objects as well as enforce that the person who uploaded the object should manage the permission of those objects. Which is the easiest way to achieve this?

- A. The root account owner should create a bucket policy which allows the IAM users to upload the object
- B. The root account owner should create the bucket policy which allows the other account owners to set the object policy of that bucket
- C. The root account should use ACL with the bucket to allow everyone to upload the object
- D. The root account should create the IAM users and provide them the permission to upload content to the bucket

Answer: C

Explanation:

Each AWS S3 bucket and object has an ACL (Access Control List. associated with it. An ACL is a list of grants identifying the grantee and the permission granted. The user can use ACLs to grant basic read/write permissions to other AWS accounts. ACLs use an Amazon S3?Vspecific XML schema. The user cannot grant permissions to other users in his account. ACLs are suitable for specific scenarios. For example, if a bucket owner allows other AWS accounts to upload objects, permissions to these objects can only be managed using the object ACL by the AWS account that owns the object.

NEW QUESTION 121

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

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

Answer: C

Explanation:

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

NEW QUESTION 125

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

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

Answer: D

Explanation:

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

NEW QUESTION 126

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

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

Answer: D

Explanation:

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

NEW QUESTION 131

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

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

Answer: C

Explanation:

The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that in-flight requests continue to be served.

NEW QUESTION 135

You are managing the AWS account of a big organization. The organization has more than 1000+ employees and they want to provide access to the various services to most of the employees. Which of the below mentioned options is the best possible solution in this case?

- A. The user should create a separate IAM user for each employee and provide access to them as per the policy
- B. The user should create an IAM role and attach STS with the rol
- C. The user should attach that role to the EC2 instance and setup AWS authentication on that server
- D. The user should create IAM groups as per the organization??s departments and add each user to the group for better access control
- E. Attach an IAM role with the organization??s authentication service to authorize each user for various AWS services

Answer: D

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The user is managing an AWS account for an organization that already has an identity system, such as the login system for the corporate network (SSO). In this case, instead of creating individual IAM users or groups for each user who need AWS access, it may be more practical to use a proxy server to translate the user identities from the organization network into the temporary AWS security credentials. This proxy server will attach an IAM role to the user after authentication.

NEW QUESTION 139

A user has configured a VPC with a new subnet. The user has created a security group. The user wants to configure that instances of the same subnet communicate with each other. How can the user configure this with the security group?

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

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user??s AWS account. AWS provides two features that the user can use to increase security in

VPC: security groups and network ACLs. Security groups work at the instance level. If the user is using the default security group it will have a rule which allows the instances to communicate with other. For a new security group the user has to specify the rule, add it to define the source as the security group itself, and select all the protocols and ports for that source.

NEW QUESTION 142

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

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

Answer: A

Explanation:

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

NEW QUESTION 143

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

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

Answer: C

Explanation:

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

NEW QUESTION 146

A user is displaying the CPU utilization, and Network in and Network out CloudWatch metrics data of a single instance on the same graph. The graph uses one Y-axis for CPU utilization and Network in and another Y-axis for Network out. Since Network in is too high, the CPU utilization data is not visible clearly on graph to the user. How can the data be viewed better on the same graph?

- A. It is not possible to show multiple metrics with the different units on the same graph
- B. Add a third Y-axis with the console to show all the data in proportion
- C. Change the axis of Network by using the Switch command from the graph
- D. Change the units of CPU utilization so it can be shown in proportion with Network

Answer: C

Explanation:

Amazon CloudWatch provides the functionality to graph the metric data generated either by the AWS services or the custom metric to make it easier for the user to analyse. It is possible to show the multiple metrics with different units on the same graph. If the graph is not plotted properly due to a difference in the unit data over two metrics, the user can change the Y-axis of one of the graph by selecting that graph and clicking on the Switch option.

NEW QUESTION 149

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

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

Answer: D

Explanation:

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

NEW QUESTION 150

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

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

Answer: C

Explanation:

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

NEW QUESTION 152

A user has created a VPC with public and private subnets using the VPC wizard. The user has not launched any instance manually and is trying to delete the VPC. What will happen in this scenario?

- A. It will not allow to delete the VPC as it has subnets with route tables
- B. It will not allow to delete the VPC since it has a running route instance
- C. It will terminate the VPC along with all the instances launched by the wizard
- D. It will not allow to delete the VPC since it has a running NAT instance

Answer: D

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance with an elastic IP. If the user is trying to delete the VPC it will not allow as the NAT instance is still running.

NEW QUESTION 154

A user has created a VPC with the public subnet. The user has created a security group for that VPC. Which of the below mentioned statements is true when a security group is created?

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

Answer: C

Explanation:

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

NEW QUESTION 156

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

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

Answer: B

Explanation:

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

NEW QUESTION 159

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 160

An organization has configured Auto Scaling with ELB. One of the instance health check returns the status as Impaired to Auto Scaling. What will Auto Scaling do in this scenario?

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

Answer: B

Explanation:

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

NEW QUESTION 164

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 166

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 170

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 172

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 175

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 177

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

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

Answer: A

Explanation:

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

NEW QUESTION 180

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 181

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

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

Answer: D

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that

subnet. A user can have 5 IP addresses per region with EC2 Classic. The user can have 5 separate IPs with VPC in the same region as it has a separate limit than EC2 Classic.

NEW QUESTION 186

A user is configuring a CloudWatch alarm on RDS to receive a notification when the CPU utilization of RDS is higher than 50%. The user has setup an alarm when there is some inactivity on RDS, such as RDS unavailability. How can the user configure this?

- A. Setup the notification when the CPU is more than 75% on RDS
- B. Setup the notification when the state is Insufficient Data
- C. Setup the notification when the CPU utilization is less than 10%
- D. It is not possible to setup the alarm on RDS

Answer: B

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 alarm has three states: Alarm, OK and Insufficient data. The Alarm will change to Insufficient Data when any of the three situations arise: when the alarm has just started, when the metric is not available or when enough data is not available for the metric to determine the alarm state. If the user wants to find that RDS is not available, he can setup to receive the notification when the state is in Insufficient data.

NEW QUESTION 191

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

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

Answer: A

Explanation:

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

NEW QUESTION 196

A user is trying to setup a security policy for ELB. The user wants ELB to meet the cipher supported by the client by configuring the server order preference in ELB security policy. Which of the below mentioned preconfigured policies supports this feature?

- A. ELBSecurity Policy-2014-01
- B. ELBSecurity Policy-2011-08
- C. ELBDefault Negotiation Policy
- D. ELBSample- OpenSSLDefault Cipher Policy

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. If the load balancer is configured to support the Server Order Preference, then the load balancer gets to select the first cipher in its list that matches any one of the ciphers in the client's list. When the user verifies the preconfigured policies supported by ELB, the policy `ELBSecurity Policy-2014-01` supports server order preference.

NEW QUESTION 199

A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling AlarmNotification (which notifies Auto Scaling for CloudWatch alarms) process for a while. What will Auto Scaling do during this period?

- A. AWS will not receive the alarms from CloudWatch
- B. AWS will receive the alarms but will not execute the Auto Scaling policy
- C. Auto Scaling will execute the policy but it will not launch the instances until the process is resumed
- D. It is not possible to suspend the AlarmNotification process

Answer: B

Explanation:

Auto Scaling performs various processes, such as Launch, Terminate Alarm Notification etc. The user can also suspend individual process. The AlarmNotification process type accepts notifications from the Amazon CloudWatch alarms that are associated with the Auto Scaling group. If the user suspends this process type, Auto Scaling will not automatically execute the scaling policies that would be triggered by the alarms.

NEW QUESTION 203

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 204

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

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

Answer: A

Explanation:

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

NEW QUESTION 208

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 209

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 212

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

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

Answer: C

Explanation:

The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private

gateway to route all traffic of the VPN subnet. If the virtual private gateway is attached with VPC and the user deletes the VPC from the console it will first detach the gateway automatically and only then delete the VPC.

NEW QUESTION 215

A user has launched an RDS MySQL DB with the Multi AZ feature. The user has scheduled the scaling of instance storage during maintenance window. What is the correct order of events during maintenance window?

Perform maintenance on standby Promote standby to primary

Perform maintenance on original primary Promote original master back as primary

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

Answer: B

Explanation:

Running MySQL on the RDS DB instance as a Multi-AZ deployment can help the user reduce the

impact of a maintenance event, as the Amazon will conduct maintenance by following the steps in the below mentioned order:

Perform maintenance on standby Promote standby to primary

Perform maintenance on original primary, which becomes the new standby.

NEW QUESTION 216

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 217

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 219

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

A. Destination: 0.0.0.0/0 and Target: i-a12345

B. Destination: 20.0.0.0/0 and Target: 80

C. Destination: 20.0.0.0/0 and Target: i-a12345

D. Destination: 20.0.0.0/24 and Target: i-a12345

Answer: A

Explanation:

A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create two route tables and attach to the subnets. The main route table will have

the entry ??Destination: 0.0.0.0/0 and Target: ia12345??, which allows all the instances in the private subnet to connect to the internet using NAT.

NEW QUESTION 221

An organization has configured Auto Scaling with ELB. There is a memory issue in the application which is causing CPU utilization to go above 90%. The higher CPU usage triggers an event for Auto Scaling as per the scaling policy. If the user wants to find the root cause inside the application without triggering a scaling activity, how can he achieve this?

A. Stop the scaling process until research is completed

B. It is not possible to find the root cause from that instance without triggering scaling

- C. Delete Auto Scaling until research is completed
- D. Suspend the scaling process until research is completed

Answer: D

Explanation:

Auto Scaling allows the user to suspend and then resume one or more of the Auto Scaling processes in the Auto Scaling group. This is very useful when the user wants to investigate a configuration problem or some other issue, such as a memory leak with the web application and then make changes to the application, without triggering the Auto Scaling process.

NEW QUESTION 222

A sys admin is planning to subscribe to the RDS event notifications. For which of the below mentioned source categories the subscription cannot be configured?

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

Answer: C

Explanation:

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

NEW QUESTION 223

A user has enabled detailed CloudWatch monitoring with the AWS Simple Notification Service. Which of the below mentioned statements helps the user understand detailed monitoring better?

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

Answer: D

Explanation:

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

NEW QUESTION 227

A user is trying to send custom metrics to CloudWatch using the PutMetricData APIs. Which of the below mentioned points should the user needs to take care while sending the data to CloudWatch?

- A. The size of a request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests
- B. The size of a request is limited to 128KB for HTTP GET requests and 64KB for HTTP POST requests
- C. The size of a request is limited to 40KB for HTTP GET requests and 8KB for HTTP POST requests
- D. The size of a request is limited to 16KB for HTTP GET requests and 80KB for HTTP POST requests

Answer: A

Explanation:

With AWS CloudWatch, the user can publish data points for a metric that share not only the same time stamp, but also the same namespace and dimensions. CloudWatch can accept multiple data points in the same PutMetricData call with the same time stamp. The only thing that the user needs to take care of is that the size of a PutMetricData request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests.

NEW QUESTION 228

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 232

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

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

Answer: D

Explanation:

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

NEW QUESTION 234

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 237

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

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

Answer: D

Explanation:

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

NEW QUESTION 242

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

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

Answer: D

Explanation:

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

NEW QUESTION 245

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 247

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

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

Answer: B

Explanation:

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

NEW QUESTION 250

A user has configured ELB with SSL using a security policy for secure negotiation between the client and load balancer. The ELB security policy supports various ciphers. Which of the below mentioned options helps identify the matching cipher at the client side to the ELB cipher list when client is requesting ELB DNS over SSL?

- A. Cipher Protocol
- B. Client Configuration Preference
- C. Server Order Preference
- D. Load Balancer Preference

Answer: C

Explanation:

Elastic Load Balancing uses a Secure Socket Layer (SSL) negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. When client is requesting ELB DNS over SSL and if the load balancer is configured to support the Server Order Preference, then the load balancer gets to select the first cipher in its list that matches any one of the ciphers in the client's list. Server Order Preference ensures that the load balancer determines which cipher is used for the SSL connection.

NEW QUESTION 252

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

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

Answer: C

Explanation:

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

NEW QUESTION 256

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 261

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

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

Answer: C

Explanation:

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

NEW QUESTION 266

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 267

A user has configured Auto Scaling with the minimum capacity as 2 and the desired capacity as 2. The user is trying to terminate one of the existing instance with the command:

```
as-terminate-instance-in-auto-scaling-group<Instance ID> --decrement-desired-capacity
```

What will Auto Scaling do in this scenario?

- A. Terminates the instance and does not launch a new instance
- B. Terminates the instance and updates the desired capacity to 1
- C. Terminates the instance and updates the desired capacity and minimum size to 1
- D. Throws an error

Answer: D

Explanation:

The Auto Scaling command `as-terminate-instance-in-auto-scaling-group <Instance ID>` will terminate the specific instance ID. The user is required to specify the parameter `--decrement-desired-capacity`. Then Auto Scaling will terminate the instance and decrease the desired capacity by 1. In this case since the minimum size is 2, Auto Scaling will not allow the desired capacity to go below 2. Thus, it will throw an error.

NEW QUESTION 270

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 274

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 276

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

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

Answer: C

Explanation:

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

NEW QUESTION 281

The CFO of a company wants to allow one of his employees to view only the AWS usage report page. Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

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

Answer: C

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow", "Action": [
        "aws-portal:ViewUsage"
      ],
      "Resource": "*"
    }
  ]
}
```

NEW QUESTION 283

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

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

Answer: D

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. AWS DynamoDB has only tables and the organization cannot make separate databases. The organization should create a table with the same name as the IAM user name and use the ARN of DynamoDB as part of the group policy. The sample policy is shown below:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": ["dynamodb:*"],
    "Resource": "arn:aws:dynamodb:region:account-number-without-hyphens:table/${aws:username}"
  }]
}
```

NEW QUESTION 285

An organization has created a Queue named ??modularqueue?? with SQS. The organization is not performing any operations such as SendMessage, ReceiveMessage, DeleteMessage, GetQueueAttributes, SetQueueAttributes, AddPermission, and RemovePermission on the queue. What can happen in this scenario?

- A. AWS SQS sends notification after 15 days for inactivity on queue
- B. AWS SQS can delete queue after 30 days without notification
- C. AWS SQS marks queue inactive after 30 days
- D. AWS SQS notifies the user after 2 weeks and deletes the queue after 3 weeks.

Answer: B

Explanation:

Amazon SQS can delete a queue without notification if one of the following actions hasn't been performed on it for 30 consecutive days: SendMessage, ReceiveMessage, DeleteMessage, GetQueueAttributes, SetQueueAttributes, AddPermission, and RemovePermission.

NEW QUESTION 288

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 293

A user is planning to scale up an application by 8 AM and scale down by 7 PM daily using Auto Scaling. What should the user do in this case?

- A. Setup the scaling policy to scale up and down based on the CloudWatch alarms
- B. The user should increase the desired capacity at 8 AM and decrease it by 7 PM manually
- C. The user should setup a batch process which launches the EC2 instance at a specific time
- D. Setup scheduled actions to scale up or down at a specific time

Answer: A

Explanation:

Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. To configure the Auto Scaling group to scale based on a schedule, the user needs to create scheduled actions. A scheduled action tells Auto Scaling to perform a scaling action at a certain time in the future.

NEW QUESTION 298

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 300

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 302

A user has moved an object to Glacier using the life cycle rules. The user requests to restore the archive after 6 months. When the restore request is completed the user accesses that archive. Which of the below mentioned statements is not true in this condition?

- A. The archive will be available as an object for the duration specified by the user during the restoration request
- B. The restored object's storage class will be RRS
- C. The user can modify the restoration period only by issuing a new restore request with the updated period
- D. The user needs to pay storage for both RRS (restore
- E. and Glacier (Archiv
- F. Rates

Answer: B

Explanation:

AWS Glacier is an archival service offered by AWS. AWS S3 provides lifecycle rules to archive and restore objects from S3 to Glacier. Once the object is archived their storage class will change to Glacier. If the user sends a request for restore, the storage class will still be Glacier for the restored object. The user will be paying for both the archived copy as well as for the restored object. The object is available only for the duration specified in the restore request and if the user wants to modify that period, he has to raise another restore request with the updated duration.

NEW QUESTION 307

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

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

Answer: D

Explanation:

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

NEW QUESTION 308

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

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

Answer: A

Explanation:

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

NEW QUESTION 309

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 311

A user has created a launch configuration for Auto Scaling where CloudWatch detailed monitoring is disabled. The user wants to now enable detailed monitoring. How can the user achieve this?

- A. Update the Launch config with CLI to set `InstanceMonitoringDisabled = false`
- B. The user should change the Auto Scaling group from the AWS console to enable detailed monitoring
- C. Update the Launch config with CLI to set `InstanceMonitoring.Enabled = true`
- D. Create a new Launch Config with detail monitoring enabled and update the Auto Scaling group

Answer: D

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 the AutoScaling 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. When the user has created a launch configuration with `InstanceMonitoring.Enabled = false` it will involve multiple steps to enable detail monitoring. The steps are:
Create a new Launch config with detailed monitoring enabled
Update the Auto Scaling group with a new launch config
Enable detail monitoring on each EC2 instance

NEW QUESTION 315

A user has launched an EC2 instance from an instance store backed AMI. The user has attached an additional instance store volume to the instance. The user wants to create an AMI from the running instance. Will the AMI have the additional instance store volume data?

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

Answer: A

Explanation:

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

NEW QUESTION 320

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

- A. The OS will shutdown but the instance will not be terminated due to protection
- B. It will terminate the instance
- C. It will not allow the user to shutdown the instance from the OS
- D. It is not possible to set the termination protection when an Instance initiated shutdown is set to Terminate

Answer: B

Explanation:

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

NEW QUESTION 322

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 324

When an EC2 instance that is backed by an S3-based AMI is terminated, what happens to the data on the root volume?

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

Answer: A

NEW QUESTION 329

In order to optimize performance for a compute cluster that requires low inter-node latency, which feature in the following list should you use?

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

Answer: B

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 330

You have a proprietary data store on-premises that must be backed up daily by dumping the data store contents to a single compressed 50GB file and sending the file to AWS. Your SLAs state that any dump file backed up within the past 7 days can be retrieved within 2 hours. Your compliance department has stated that all data must be held indefinitely. The time required to restore the data store from a backup is approximately 1 hour. Your on-premise network connection is capable of sustaining 1gbps to AWS.

Which backup methods to AWS would be most cost-effective while still meeting all of your requirements?

- A. Send the daily backup files to Glacier immediately after being generated
- B. Transfer the daily backup files to an EBS volume in AWS and take daily snapshots of the volume
- C. Transfer the daily backup files to S3 and use appropriate bucket lifecycle policies to send to Glacier
- D. Host the backup files on a Storage Gateway with Gateway-Cached Volumes and take daily snapshots

Answer: D

Explanation:

Reference:

<http://aws.amazon.com/storagegateway/faqs/>

NEW QUESTION 332

.....

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](#)