

Exam Questions AWS-SysOps

Amazon AWS Certified SysOps Administrator - Associate

<https://www.2passeasy.com/dumps/AWS-SysOps/>



NEW QUESTION 1

- (Topic 1)

A media company produces new video files on-premises every day with a total size of around 100GBS after compression All files have a size of 1 -2 GB and need to be uploaded to Amazon S3 every night in a fixed time window between 3am and 5am Current upload takes almost 3 hours, although less than half of the available bandwidth is used.

What step(s) would ensure that the file uploads are able to complete in the allotted time window?

- A. Increase your network bandwidth to provide faster throughput to S3
- B. Upload the files in parallel to S3
- C. Pack all files into a single archive, upload it to S3, then extract the files in AWS
- D. Use AWS Import/Export to transfer the video files

Answer: B

Explanation:

Reference:

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

NEW QUESTION 2

- (Topic 1)

Your EC2-Based Multi-tier application includes a monitoring instance that periodically makes application -level read only requests of various application components and if any of those fail more than three times 30 seconds calls CloudWatch to fire an alarm, and the alarm notifies your operations team by email and SMS of a possible application health problem. However, you also need to watch the watcher -the monitoring instance itself - and be notified if it becomes unhealthy.

Which of the following is a simple way to achieve that goal?

- A. Run another monitoring instance that pings the monitoring instance and fires a CloudWatch alarm that notifies your operations team should the primary monitoring instance become unhealthy
- B. Set a CloudWatch alarm based on EC2 system and instance status checks and have the alarm notify your operations team of any detected problem with the monitoring instance
- C. Set a CloudWatch alarm based on the CPU utilization of the monitoring instance and have the alarm notify your operations team if the CPU usage exceeds 50% for more than one minute: then have your monitoring application go into a CPU-bound loop should it detect any application problem
- D. Have the monitoring instances post messages to an SQS queue and then dequeue those messages on another instance should the queue cease to have new messages, the second instance should first terminate the original monitoring instance start another backup monitoring instance and assume the role of the previous monitoring instance and begin adding messages to the SQS queue

Answer: D

NEW QUESTION 3

- (Topic 1)

How can the domain's zone apex for example "myzoneapexdomain.com" be pointed towards an Elastic Load Balancer?

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 4

- (Topic 1)

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

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

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

Answer: C

Explanation:

Explanation: Reference:

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

NEW QUESTION 5

- (Topic 1)

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

Choose 3 answers

- A. Gather evidence of your IT operational controls
- B. Request and obtain applicable third-party audited AWS compliance reports and certifications

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

Answer: ABD

NEW QUESTION 6

- (Topic 1)

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

Which two options will address this issue? Choose 2 answers.

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

Answer: AB

NEW QUESTION 7

- (Topic 1)

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

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 8

- (Topic 1)

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

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

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 9

- (Topic 1)

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: AD

NEW QUESTION 10

- (Topic 1)

You have decided to change the Instance type for instances running in your application tier that are using Auto Scaling.

In which area below would you change the instance type definition?

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

Answer: A

Explanation:

Reference:

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/WhatIsAutoScaling.html>

NEW QUESTION 10

- (Topic 1)

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

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

Answer: C

NEW QUESTION 14

- (Topic 1)

Your entire AWS infrastructure lives inside of one Amazon VPC You have an Infrastructure monitoring application running on an Amazon instance in Availability Zone (AZ) A of the region, and another application instance running in AZ B. The monitoring application needs to make use of ICMP ping to confirm network reachability of the instance hosting the application.

Can you configure the security groups for these instances to only allow the ICMP ping to pass from the monitoring instance to the application instance and nothing else" If so how?

- A. No Two instances in two different AZ's can't talk directly to each other via ICMP ping as that protocol is not allowed across subnet (iebroadcast) boundaries
- B. Yes Both the monitoring instance and the application instance have to be a part of the same security group, and that security group needs to allow inbound ICMP
- C. Yes, The security group for the monitoring instance needs to allow outbound ICMP and the application instance's security group needs to allow Inbound ICMP
- D. Yes, Both the monitoring instance's security group and the application instance's security group need to allow both inbound and outbound ICMP ping packets since ICMP is not a connection-oriented protocol

Answer: D

NEW QUESTION 17

- (Topic 1)

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

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

Answer: B

NEW QUESTION 21

- (Topic 1)

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

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

Answer: ACD

Explanation:

Reference:

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

NEW QUESTION 23

- (Topic 1)

You have a web application leveraging an Elastic Load Balancer (ELB) In front of the web servers deployed using an Auto Scaling Group Your database is running on Relational

Database Service (RDS) The application serves out technical articles and responses to them in general there are more views of an article than there are responses to the article. On occasion, an article on the site becomes extremely popular resulting in significant traffic Increases that causes the site to go down. What could you do to help alleviate the pressure on the infrastructure while maintaining availability during these events?

Choose 3 answers

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

Answer: ACE

NEW QUESTION 26

- (Topic 1)

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: B

Explanation:

Reference:

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deployment.html>

NEW QUESTION 29

- (Topic 1)

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

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

Choose 2 answers.

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

Answer: AB

NEW QUESTION 33

- (Topic 1)

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

How can you accomplish this?

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

Answer: C

NEW QUESTION 37

- (Topic 1)

You are currently hosting multiple applications in a VPC and have logged numerous port scans coming in from a specific IP address block. Your security team has requested that all access from the offending IP address block be denied for the next 24 hours.

Which of the following is the best method to quickly and temporarily deny access from the specified IP address block?

- A. Create an AD policy to modify Windows Firewall settings on all hosts in the VPC to deny access from the IP address block.
- B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP address block.
- C. Add a rule to all of the VPC's Security Groups to deny access from the IP address block.
- D. Modify the Windows Firewall settings on all Amazon Machine Images (AMIs) that your organization uses in that VPC to deny access from the IP address block.

Answer: B

Explanation:

Reference:

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

NEW QUESTION 39

- (Topic 1)

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: D

NEW QUESTION 41

- (Topic 1)

Your application currently leverages AWS Auto Scaling to grow and shrink as load increases/ decreases and has been performing well. Your marketing team expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks. Your forecast for the approximate number of Amazon EC2 instances necessary to meet the peak demand is 175.

What should you do to avoid potential service disruptions during the ramp up in traffic?

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

Answer: D

NEW QUESTION 45

- (Topic 1)

Which of the following are characteristics of Amazon VPC subnets?

Choose 2 answers

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

Answer: CE

NEW QUESTION 49

- (Topic 1)

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 saved as an EBS volume
- B. Data is automatically saved as an EBS 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 54

- (Topic 1)

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

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 56

- (Topic 1)

Your company is moving towards tracking web page users with a small tracking

image loaded on each page. Currently you are serving this image out of US-East, but are starting to get concerned about the time it takes to load the image for users on the west coast.

What are the two best ways to speed up serving this image?

Choose 2 answers

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

Answer: AD

NEW QUESTION 58

- (Topic 1)

You have started a new job and are reviewing your company's infrastructure on AWS. You notice one web application where they have an Elastic Load Balancer (ELB) in front of web instances in an Auto Scaling Group. When you check the metrics for the ELB in CloudWatch, you see four healthy instances in Availability Zone (AZ) A and zero in AZ B. There are zero unhealthy instances.

What do you need to fix to balance the instances across AZs?

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

Answer: B

NEW QUESTION 60

- (Topic 1)

Which two AWS services provide out-of-the-box user configurable automatic backup-as-a-service and backup rotation options?

Choose 2 answers

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

Answer: BD

NEW QUESTION 65

- (Topic 1)

Which services allow the customer to retain full administrative privileges of the underlying EC2 instances?

Choose 2 answers

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

Answer: AC

NEW QUESTION 68

- (Topic 1)

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

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

Answer: CD

NEW QUESTION 72

- (Topic 1)

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 73

- (Topic 1)

You have been asked to automate many routine systems administrator backup and recovery activities. Your current plan is to leverage AWS-managed solutions as much as possible and automate the rest with the AWS CLI and scripts.

Which task would be best accomplished with a script?

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

Answer: A

NEW QUESTION 77

- (Topic 1)

What are characteristics of Amazon S3? Choose 2 answers

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

Answer: AD

NEW QUESTION 79

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 80

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 82

- (Topic 2)

A user is planning to use AWS Cloudformation. Which of the below mentioned functionalities does not help him to correctly understand Cloudformation?

- A. Cloudformation follows the DevOps model for the creation of Dev & Test
- B. AWS Cloudformation does not charge the user for its service but only charges for the AWS resources created with it
- C. Cloudformation works with a wide variety of AWS services, such as EC2, EBS, VPC, IAM, S3, RDS, ELB, etc
- D. CloudFormation provides a set of application bootstrapping scripts which enables the user to install Software

Answer: A

Explanation:

AWS Cloudformation is an application management tool which provides application modelling, deployment, configuration, management and related activities. It supports a wide variety of AWS services, such as EC2, EBS, AS, ELB, RDS, VPC, etc. It also provides application bootstrapping scripts which enable the user to install software packages or create folders. It is free of the cost and only charges the user for the services created with it. The only challenge is that it does not follow any model, such as DevOps; instead customers can define templates and use them to provision and manage the AWS resources in an orderly way.

NEW QUESTION 86

- (Topic 2)

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

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

Answer: B

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic

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

NEW QUESTION 89

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 90

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 95

- (Topic 2)

A user has setup an EBS backed instance and a CloudWatch alarm when the CPU utilization is more than 65%. The user has setup the alarm to watch it for 5 periods of 5 minutes each. The CPU utilization is 60% between 9 AM to 6 PM. The user has stopped the EC2 instance for 15 minutes between 11 AM to 11:15 AM. What will be the status of the alarm at 11:30 AM?

- A. Alarm
- B. OK
- C. Insufficient Data
- D. Error

Answer: B

Explanation:

Amazon CloudWatch alarm watches a single metric over a time period 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 state of the alarm will be OK for the whole day. When the user stops the instance for three periods the alarm may not receive the data

NEW QUESTION 100

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 103

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 107

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 110

- (Topic 2)

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

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

Answer: B

Explanation:

Amazon CloudWatch is basically a metrics repository. Either the user can send the custom data or an AWS product can put metrics into the repository, and the user can retrieve the statistics based on those metrics. The statistics are metric data aggregations over specified periods of time. Aggregations are made using the namespace, metric name, dimensions, and the data point unit of measure, within the time period that is specified by the user. To aggregate the data across instances launched with AMI, the user should select the AMI ID under EC2 metrics and select the aggregate average to view the data.

NEW QUESTION 114

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 119

- (Topic 2)

A user has created a queue named "myqueue" in US-East region with AWS SQS. The user's AWS account ID is 123456789012. If the user wants to perform

some action on this queue, which of the below Queue URL should he use?

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

Answer: A

Explanation:

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

NEW QUESTION 121

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 126

- (Topic 2)

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 131

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 134

- (Topic 2)

A user has launched an ELB which has 5 instances registered with it. The user deletes the ELB by mistake. What will happen to the instances?

- A. ELB will ask the user whether to delete the instances or not
- B. Instances will be terminated
- C. ELB cannot be deleted if it has running instances registered with it
- D. Instances will keep running

Answer: D

Explanation:

When the user deletes the Elastic Load Balancer, all the registered instances will be deregistered. However, they will continue to run. The user will incur charges if he does not take any action on those instances.

NEW QUESTION 135

- (Topic 2)

An organization has setup consolidated billing with 3 different AWS accounts. Which of the below mentioned advantages will organization receive in terms of the AWS pricing?

- A. The consolidated billing does not bring any cost advantage for the organization
- B. All AWS accounts will be charged for S3 storage by combining the total storage of each account
- C. The EC2 instances of each account will receive a total of 750*3 micro instance hours free
- D. The free usage tier for all the 3 accounts will be 3 years and not a single year

Answer: B

Explanation:

AWS consolidated billing enables the organization to consolidate payments for multiple Amazon Web Services (AWS) accounts within a single organization by making a single paying account. For billing purposes, AWS treats all the accounts on the consolidated bill as one account. Some services, such as Amazon EC2 and Amazon S3 have volume pricing tiers across certain usage dimensions that give the user lower prices when he uses the service more.

NEW QUESTION 139

- (Topic 2)

A user has setup connection draining with ELB to allow in-flight requests to continue while the instance is being deregistered through Auto Scaling. If the user has not specified the draining time, how long will ELB allow 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 144

- (Topic 2)

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 149

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 150

- (Topic 2)

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 152

- (Topic 2)

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: B

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.

NEW QUESTION 155

- (Topic 2)

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-specific 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 156

- (Topic 2)

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 160

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 162

- (Topic 2)

A user is trying to configure the CloudWatch billing alarm. Which of the below mentioned steps should be performed by the user for the first time alarm creation in the AWS Account Management section?

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

Answer: B

Explanation:

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

NEW QUESTION 167

- (Topic 2)

A user wants to make so that whenever the CPU utilization of the AWS EC2 instance is above 90%, the redlight of his bedroom turns on. Which of the below mentioned AWS services is helpful for this purpose?

- A. AWS CloudWatch + AWS SES
- B. AWS CloudWatch + AWS SNS
- C. Non
- D. It is not possible to configure the light with the AWS infrastructure services
- E. AWS CloudWatch and a dedicated software turning on the light

Answer: B

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 configure some sensor devices at his home which receives data on the HTTP end point (REST calls) and turn on the red light. The user can configure the CloudWatch alarm to send a notification to the AWS SNS HTTP end point (the sensor device) and it will turn the light red when there is an alarm condition.

NEW QUESTION 171

- (Topic 2)

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

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

Answer: D

Explanation:

A user can reboot an EC2 instance using the AWS console, the Amazon EC2 CLI or the Amazon EC2 API. Rebooting an instance is equivalent to rebooting an operating system. However, it is recommended that the user use the Amazon EC2 to reboot the instance instead of running the operating system reboot command

from the instance. The instance remains on the same host computer and maintains its public DNS name, private IP address, and any data on its instance store volumes. It typically takes a few minutes for the reboot to complete, but the time it takes to reboot depends on the instance configuration.

NEW QUESTION 175

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 180

- (Topic 2)

An organization (Account ID 123412341234. has attached the below mentioned IAM policy to a user. What does this policy statement entitle the user to perform?

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

- A. 0
- B. 0
- C. 0
- D. 0

Answer: A

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (Account ID 123412341234. wants some of their users to manage keys (access and secret access keys. of all IAM users, the organization should set the below mentioned policy which entitles the IAM user to modify keys of all IAM users with CLI, SDK or API.

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

NEW QUESTION 183

- (Topic 2)

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

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

Answer: C

Explanation:

Amazon CloudWatch aggregates statistics according to the period length that the user has specified while getting data from CloudWatch. The user can publish as many data points as he wants with the same or 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 187

- (Topic 2)

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 192

- (Topic 2)

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. 0
- C. 5
- D. 2

Answer: C

NEW QUESTION 197

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 201

- (Topic 2)

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 205

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 210

- (Topic 2)

A user has setup Auto Scaling with ELB on the EC2 instances. The user wants to configure that whenever the CPU utilization is below 10%, Auto Scaling should remove one instance. How can the user configure this?

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

Answer: D

Explanation:

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

NEW QUESTION 211

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 212

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 213

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 218

- (Topic 2)

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 222

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 226

- (Topic 2)

A system admin is planning to setup event notifications on RDS. Which of the below mentioned services will help the admin setup notifications?

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

Answer: D

Explanation:

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

NEW QUESTION 229

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 234

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 236

- (Topic 2)

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 240

- (Topic 3)

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 244

- (Topic 3)

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 249

- (Topic 3)

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 254

- (Topic 3)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 259

- (Topic 3)

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 260

- (Topic 3)

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 261

- (Topic 3)

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

happen?

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

Answer: D

Explanation:

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

NEW QUESTION 264

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 269

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 273

- (Topic 3)

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

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

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

Answer: C

NEW QUESTION 276

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 277

- (Topic 3)

You have a business-to-business web application running in a VPC consisting of an Elastic Load Balancer (ELB), web servers, application servers and a database. Your web application should only accept traffic from pre-defined customer IP addresses.

Which two options meet this security requirement? Choose 2 answers A. Configure web server VPC security groups to allow traffic from your customers' IPs

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

Answer: AB

NEW QUESTION 282

- (Topic 3)

An organization has created one IAM user and applied the below mentioned policy to the user. What entitlements do the IAM users avail with this policy?

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:Describe*",
      "Resource": "*"
    },
    {
      "Effect": "Allow"
```

```

"Action": [
"cloudwatch:ListMetrics",
"cloudwatch:GetMetricStatistics",
"cloudwatch:Describe*"
],
"Resource": "*"
},
{
"Effect": "Allow",
"Action": "autoscaling:Describe*",
"Resource": "*"
}
]
}

```

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

Answer: D

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If an organization wants to setup read only access to EC2 for a particular user, they should mention the action in the IAM policy which entitles the user for Describe rights for EC2, CloudWatch, Auto Scaling and ELB. In the policy shown below, the user will have read only access for EC2 and EBS, CloudWatch and Auto Scaling. Since ELB is not mentioned as a part of the list, the user will not have access to ELB.

```

{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "ec2:Describe*",
"Resource": "*"
},
{
"Effect": "Allow",
"Action": [
"cloudwatch:ListMetrics",
"cloudwatch:GetMetricStatistics",
"cloudwatch:Describe*"
],
"Resource": "*"
},
{
"Effect": "Allow",
"Action": "autoscaling:Describe*",
"Resource": "*"
}
]
}

```

NEW QUESTION 284

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 286

- (Topic 3)

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

- A. Add a new database layer and then add recipes to the deploy actions of the database and App Server layer
- B. Use OpsWorks' "Clone Stack" feature to create a second RDS stack in another Availability Zone for redundancy in the event of a failure in the Primary A
- C. To switch to the secondary RDS instance, set the `[:database]` attributes to values that are appropriate for your server which you can do by using custom JSO
- D. The variables that characterize the RDS database connection—host, user, and so on—are set using the corresponding values from the deploy JSON's

[[:deploy]][[:app_name]][[:database]] attribute

E. Cookbook attributes are stored in a repository, so OpsWorks requires that the "password": "your_password" attribute for the RDS instance must be encrypted using at least a 256-bit key

F. Set up the connection between the app server and the RDS layer by using a custom recipe

G. The recipe configures the app server as required, typically by creating a configuration file

H. The recipe gets the connection data such as the host and database name from a set of attributes in the stack configuration and deployment JSON that AWS OpsWorks installs on every instance

Answer: BCE

NEW QUESTION 291

- (Topic 3)

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 296

- (Topic 3)

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

A. Create a bucket policy and apply it to the bucket

B. Create a NACL and attach it to the VPC of the bucket

C. Create an ACL and apply it to all objects in the bucket

D. Modify the IAM policies of any users that would access the bucket

Answer: A

Explanation:

Reference:

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

NEW QUESTION 298

- (Topic 3)

In AWS, which security aspects are the customer's responsibility? Choose 4 answers

A. Controlling physical access to compute resources

B. Patch management on the EC2 instances operating system

C. Encryption of EBS (Elastic Block Storage) volumes

D. Life-cycle management of IAM credentials

E. Decommissioning storage devices

F. Security Group and ACL (Access Control List) settings

Answer: BCEF

NEW QUESTION 303

- (Topic 3)

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 304

- (Topic 3)

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

understand detailed monitoring better?

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

Answer: D

Explanation:

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

NEW QUESTION 309

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 314

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 318

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 319

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 321

- (Topic 3)

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 322

- (Topic 3)

A user has configured Auto Scaling with 3 instances. The user had created a new AMI after updating one of the instances. If the user wants to terminate two specific instances to ensure that Auto Scaling launches an instances with the new launch configuration, which command should he run?

- A. `as-delete-instance-in-auto-scaling-group <Instance ID> --no-decrement-desired-capacity`
- B. `as-terminate-instance-in-auto-scaling-group <Instance ID> --update-desired-capacity`
- C. `as-terminate-instance-in-auto-scaling-group <Instance ID> --decrement-desired-capacity`
- D. `as-terminate-instance-in-auto-scaling-group <Instance ID> --no-decrement-desired-capacity`

Answer: D

Explanation:

The Auto Scaling command `as-terminate-instance-in-auto-scaling-group <Instance ID>` will terminate the specific instance ID. The user is required to specify the parameter `--no-decrement-desired-capacity` to ensure that it launches a new instance from the launch config after terminating the instance. If the user specifies the parameter `--decrement-desired-capacity` then Auto Scaling will terminate the instance and decrease the desired capacity by 1.

NEW QUESTION 324

- (Topic 3)

An organization (Account ID 123412341234, has attached the below mentioned IAM policy to a user. What does this policy statement entitle the user to perform?

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowUsersAllActionsForCredentials",
    "Effect": "Allow",
    "Action": [
      "iam:*LoginProfile",
      "iam:*AccessKey*",
      "iam:*SigningCertificate*"
    ],
    "Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
  }]
}
```

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

Answer: D

Explanation:

WS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (Account ID 123412341234, wants some of their users to manage credentials (access keys, password, and sign in certificates, of all IAM users, they should set an applicable policy to that user or group of users. The below mentioned policy allows the IAM user to modify the credentials of all IAM user's using only CLI, SDK or APIs. The user cannot use the AWS console for this activity since he does not have list permission for the IAM users.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowUsersAllActionsForCredentials",
    "Effect": "Allow"
```

```
"Action": [  
  "iam:*LoginProfile",  
  "iam:*AccessKey*",  
  "iam:*SigningCertificate*" ]  
},  
"Resource": ["arn:aws:iam::123412341234:user/${aws:username}"]  
Amazon AWS-SysOps : Practice Test  
}}}
```

NEW QUESTION 328

- (Topic 3)

The compliance department within your multi-national organization requires that all data for your customers that reside in the European Union (EU) must not leave the EU and also

data for customers that reside in the US must not leave the US without explicit authorization.

What must you do to comply with this requirement for a web based profile management application running on EC2?

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

Answer: C

NEW QUESTION 330

- (Topic 3)

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 332

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 333

- (Topic 3)

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

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

What will Auto Scaling do in this scenario?

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

Answer: D

Explanation:

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

NEW QUESTION 338

- (Topic 3)

A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services provides detailed monitoring with CloudWatch without charging the user extra?

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

Answer: B

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. Services, such as RDS, ELB, OpsWorks, and Route 53 can provide the monitoring data every minute without charging the user.

NEW QUESTION 339

- (Topic 3)

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

How can the user achieve this?

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

Answer: A

Explanation:

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

NEW QUESTION 343

- (Topic 3)

When you put objects in Amazon S3, what is the indication that an object was successfully stored?

- A. Each S3 account has a special bucket named `_s3_log`
- B. Success codes are written to this bucket with a timestamp and checksum
- C. A success code is inserted into the S3 object metadata
- D. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful
- E. Amazon S3 is engineered for 99.999999999% durability
- F. Therefore there is no need to confirm that data was inserted

Answer: B

Explanation:

Reference:

<http://docs.aws.amazon.com/AmazonS3/latest/API/RESTObjectPUT.html>

NEW QUESTION 346

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 349

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 352

- (Topic 3)

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

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

Answer: C

Explanation:

When the user has launched an EC2 instance from an instance store backed AMI and the admin team wants to create an AMI from it, the user needs to setup the AWS CLI or the API tools first. Once the tool is setup the user will need the following credentials:

- AWS account ID;
- AWS access and secret access key;
- X.509 certificate with private key.

NEW QUESTION 353

- (Topic 3)

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 357

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 359

- (Topic 3)

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

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

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

Answer: B

NEW QUESTION 364

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 369

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 371

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 374

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 379

- (Topic 3)

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 384

- (Topic 3)

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 385

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 389

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 392

- (Topic 3)

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 397

- (Topic 3)

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

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

Answer: B

NEW QUESTION 401

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-SysOps Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-SysOps Product From:

<https://www.2passeasy.com/dumps/AWS-SysOps/>

Money Back Guarantee

AWS-SysOps Practice Exam Features:

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