



Amazon

Exam Questions AWS-SysOps

Amazon AWS Certified SysOps Administrator - Associate

NEW QUESTION 1

- (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 Provisioned IOPS EBS volumes in a RAID 0 configuration

Answer: C

Explanation:

Explanation: Reference:

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

NEW QUESTION 2

- (Topic 1)

You need to design a VPC for a web-application consisting of an Elastic Load Balancer (ELB), a fleet of web/application servers, and an RDS database. The entire infrastructure must be distributed over 2 availability zones.

Which VPC configuration works while assuring the database is not available from the Internet?

- A. One public subnet for ELB, one public subnet for the web-servers, and one private subnet for the database
- B. One public subnet for ELB, two private subnets for the web-servers, two private subnets for RDS
- C. Two public subnets for ELB, two private subnets for the web-servers, and two private subnets for RDS
- D. Two public subnets for ELB, two public subnets for the web-servers, and two public subnets for RDS

Answer: A

NEW QUESTION 3

- (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 ELB
- E. RDS and EC2 instances

Answer: B

NEW QUESTION 4

- (Topic 1)

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

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

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

Answer: B

NEW QUESTION 5

- (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 IAM 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 6

- (Topic 1)

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

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 7

- (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 8

- (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 9

- (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 Redshift

Answer: BD

NEW QUESTION 10

- (Topic 1)

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

What two actions could you take to rectify this?

Choose 2 answers

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

Answer: BD

NEW QUESTION 10

- (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 14

- (Topic 1)

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

Which configuration would achieve that goal?

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 19

- (Topic 1)

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

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

Answer: CD

NEW QUESTION 21

- (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 26

- (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 30

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 32

- (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 34

- (Topic 2)

A user wants to disable connection draining on an existing ELB. Which of the below mentioned statements helps the user disable connection draining on the ELB?

- A. The user can only disable connection draining from CLI
- B. It is not possible to disable the connection draining feature once enabled
- C. The user can disable the connection draining feature from EC2 -> ELB console or from CLI
- D. The user needs to stop all instances before disabling connection draining

Answer: C

Explanation:

The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that inflight requests continue to be served. The user can enable or disable connection draining from the AWS EC2 console -> ELB or using CLI.

NEW QUESTION 35

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 38

- (Topic 2)

A user has created an ELB with the availability zone US-East-1A. The user wants to add more zones to ELB to achieve High Availability. How can the user add more zones to the existing ELB?

- A. It is not possible to add more zones to the existing ELB
- B. The only option is to launch instances in different zones and add to ELB
- C. The user should stop the ELB and add zones and instances as required
- D. The user can add zones on the fly from the AWS console

Answer: D

Explanation:

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

NEW QUESTION 39

- (Topic 2)

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

- A. AWS Simple Notification Service
- B. AWS Simple Workflow

- C. AWS Simple Queue Service
- D. AWS Simple Query Service

Answer: C

Explanation:

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

NEW QUESTION 41

- (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 46

- (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. The user is trying to create the private subnet with CIDR 20.0.0.128/25. Which of the below mentioned statements is true in this scenario?

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

Answer: B

Explanation:

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

NEW QUESTION 50

- (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 52

- (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 inflight 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 57

- (Topic 2)

A user has created an S3 bucket which is not publicly accessible. The bucket is having thirty objects which are also private. If the user wants to make the objects public, how can he configure this with minimal efforts?

- A. The user should select all objects from the console and apply a single policy to mark them public
- B. The user can write a program which programmatically makes all objects public using S3 SDK
- C. Set the AWS bucket policy which marks all objects as public
- D. Make the bucket ACL as public so it will also mark all objects as public

Answer: C

Explanation:

A system admin can grant permission of the S3 objects or buckets to any user or make the objects public using the bucket policy and user policy. Both use the JSON-based access policy language. Generally if the user is defining the ACL on the bucket, the objects in the bucket do not inherit it and vice versa. The bucket policy can be defined at the bucket level which allows the objects as well as the bucket to be public with a single policy applied to that bucket.

NEW QUESTION 59

- (Topic 2)

A customer is using AWS for Dev and Test. The customer wants to setup the Dev environment with Cloudformation. Which of the below mentioned steps are not required while using Cloudformation?

- A. Create a stack
- B. Configure a service
- C. Create and upload the template
- D. Provide the parameters configured as part of the template

Answer: B

Explanation:

AWS Cloudformation is an application management tool which provides application modelling, deployment, configuration, management and related activities. AWS CloudFormation introduces two concepts: the template and the stack. The template is a JSON-format, text-based file that describes all the AWS resources required to deploy and run an application. The stack is a collection of AWS resources which are created and managed as a single unit when AWS CloudFormation instantiates a template. While creating a stack, the user uploads the template and provides the data for the parameters if required.

NEW QUESTION 63

- (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 64

- (Topic 2)

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

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

Answer: C

Explanation:

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. In a Multi-AZ deployment, Amazon RDS automatically

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

NEW QUESTION 68

- (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 69

- (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 72

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 74

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 76

- (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 similartime 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 77

- (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 81

- (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 86

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 87

- (Topic 2)

A user is launching an instance. He is on the "Tag the instance" screen. Which of the below mentioned information will not help the user understand the functionality of an AWS tag?

- A. Each tag will have a key and value
- B. The user can apply tags to the S3 bucket
- C. The maximum value of the tag key length is 64 unicode characters
- D. AWS tags are used to find the cost distribution of various resources

Answer: C

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. The maximum size of a tag key is 128 unicode characters.

NEW QUESTION 88

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 91

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 96

- (Topic 2)

An organization wants to move to Cloud. They are looking for a secure encrypted database storage option. Which of the below mentioned AWS functionalities helps them to achieve this?

- A. AWS MFA with EBS
- B. AWS EBS encryption
- C. Multi-tier encryption with Redshift
- D. AWS S3 server side storage

Answer: B

Explanation:

AWS EBS supports encryption of the volume while creating new volumes. 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 EBS will be encrypted. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between the EC2 instances and EBS storage. EBS encryption is based on the AES-256 cryptographic algorithm, which is the industry standard

NEW QUESTION 98

- (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 101

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 104

- (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 105

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 108

- (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 113

- (Topic 3)

A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling terminate process only for a while. What will happen to the availability zone rebalancing process (AZRebalance) during this period?

- A. Auto Scaling will not launch or terminate any instances
- B. Auto Scaling will allow the instances to grow more than the maximum size
- C. Auto Scaling will keep launching instances till the maximum instance size
- D. It is not possible to suspend the terminate process while keeping the launch active

Answer: B

Explanation:

Auto Scaling performs various processes, such as Launch, Terminate, Availability Zone Rebalance (AZRebalance) etc. The AZRebalance process type seeks to maintain a balanced number of instances across Availability Zones within a region. If the user suspends the Terminate process, the AZRebalance process can cause the Auto Scaling group to grow up to ten percent larger than the maximum size. This is because Auto Scaling allows groups to temporarily grow larger than the maximum size during rebalancing activities. If Auto Scaling cannot terminate instances, the Auto Scaling group could remain up to ten percent larger than the maximum size until the user resumes the Terminate process type.

NEW QUESTION 114

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 117

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 118

- (Topic 3)

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

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

Answer: A

NEW QUESTION 121

- (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 125

- (Topic 3)

A system admin is planning to encrypt all objects being uploaded to S3 from an application. The system admin does not want to implement his own encryption algorithm; instead he is planning to use server side encryption by supplying his own key (SSE-C.. Which parameter is not required while making a call for SSE-C?

- A. x-amz-server-side-encryption-customer-key-AES-256
- B. x-amz-server-side-encryption-customer-key
- C. x-amz-server-side-encryption-customer-algorithm
- D. x-amz-server-side-encryption-customer-key-MD5

Answer: A

Explanation:

AWS S3 supports client side or server side encryption to encrypt all data at rest. The server side encryption can either have the S3 supplied AES-256 encryption key or the user can send the key along with each API call to supply his own encryption key (SSE-C.. When the user is supplying his own encryption key, the user has to send the below mentioned parameters as a part of the API calls: x-amz-server-side-encryption-customer-algorithm: Specifies the encryption algorithm x-amz-server-side-encryption-customer-key: To provide the base64-encoded encryption key x-amz-server-side-encryption-customer-key-MD5: To provide the base64-encoded 128-bit MD5 digest of the encryption key

NEW QUESTION 129

- (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 134

- (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 139

- (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 144

- (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 147

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 150

- (Topic 3)

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

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 152

- (Topic 3)

A user is collecting 1000 records per second. The user wants to send the data to CloudWatch using the custom namespace. Which of the below mentioned options is recommended for this activity?

- A. Aggregate the data with statistics, such as Min, max, Average, Sum and Sample data and send the data to CloudWatch
- B. Send all the data values to CloudWatch in a single command by separating them with a comm
- C. CloudWatch will parse automatically
- D. Create one csv file of all the data and send a single file to CloudWatch
- E. It is not possible to send all the data in one cal
- F. Thus, it should be sent one by on
- G. CloudWatch will aggregate the data automatically

Answer: A

Explanation:

AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. The user can publish data to CloudWatch as single data points or as an aggregated set of data points called a statistic set using the command `put-metric-data`. It is recommended that when the user is having multiple data points per minute, he should aggregate the data so that it will minimize the number of calls to `put-metric-data`. In this case it will be single call to CloudWatch instead of 1000 calls if the data is aggregated.

NEW QUESTION 157

- (Topic 3)

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

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

Answer: AB

NEW QUESTION 158

- (Topic 3)

A user runs the command `dd if=/dev/zero of=/dev/xvdfbs=1M` on a fresh blank EBS volume attached to a Linux instance. Which of the below mentioned activities is the user performing with the command given above?

- A. Creating a file system on the EBS volume
- B. Mounting the device to the instance
- C. Pre warming the EBS volume
- D. Formatting the EBS volume

Answer: C

Explanation:

When the user creates a new EBS volume and is trying to access it for the first time it will encounter reduced IOPS due to wiping or initiating of the block storage. To avoid this as well as achieve the best performance it is required to pre warm the EBS volume. For a blank volume attached with a Linux OS, the `dd` command is used to write to all the blocks on the device. In the command `dd if=/dev/zero of=/dev/xvdfbs=1M` the parameter `if=import file` should be set to one of the Linux virtual devices, such as `/dev/zero`. The `of=output file` parameter should be set to the drive that the user wishes to warm. The `bs` parameter sets the block size of the write operation; for optimal performance, this should be set to 1 MB.

NEW QUESTION 161

- (Topic 3)

A user is planning to use AWS services for his web application. If the user is trying to set up his own billing management system for AWS, how can he configure it?

- A. Set up programmatic billing acces
- B. Download and parse the bill as per the requirement
- C. It is not possible for the user to create his own billing management service with AWS

- D. Enable the AWS CloudWatch alarm which will provide APIs to download the alarm data
- E. Use AWS billing APIs to download the usage report of each service from the AWS billing console

Answer: A

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. AWS will upload the bill to the bucket every few hours and the user can download the bill CSV from the bucket, parse it and create a billing system as per the requirement.

NEW QUESTION 165

- (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 167

- (Topic 3)

A user is using CloudFormation to launch an EC2 instance and then configure an application after the instance is launched. The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly. How can the user configure this?

- A. It is not possible that the stack creation will wait until one service is created and launched
- B. The user can use the HoldCondition resource to wait for the creation of the other dependent resources
- C. The user can use the DependentCondition resource to hold the creation of the other dependent resources
- D. The user can use the WaitCondition resource to hold the creation of the other dependent resources

Answer: D

Explanation:

AWS CloudFormation is an application management tool which provides application modelling, deployment, configuration, management and related activities. AWS CloudFormation provides a WaitCondition resource which acts as a barrier and blocks the creation of other resources until a completion signal is received from an external source, such as a user application or management system.

NEW QUESTION 172

- (Topic 3)

A user has created an EBS volume of 10 GB and attached it to a running instance. The user is trying to access EBS for first time. Which of the below mentioned options is the correct statement with respect to a first time EBS access?

- A. The volume will show a size of 8 GB
- B. The volume will show a loss of the IOPS performance the first time
- C. The volume will be blank
- D. If the EBS is mounted it will ask the user to create a file system

Answer: B

Explanation:

A user can create an EBS volume either from a snapshot or as a blank volume. If the volume is from a snapshot it will not be blank. The volume shows the right size only as long as it is mounted. This shows that the file system is created. When the user is accessing the volume the AWS EBS will wipe out the block storage or instantiate from the snapshot. Thus, the volume will show a loss of IOPS. It is recommended that the user should pre warm the EBS before use to achieve better IO.

NEW QUESTION 176

- (Topic 3)

A user is creating a CloudFormation stack. Which of the below mentioned limitations does not hold true for CloudFormation?

- A. One account by default is limited to 100 templates
- B. The user can use 60 parameters and 60 outputs in a single template
- C. The template, parameter, output, and resource description fields are limited to 4096 characters
- D. One account by default is limited to 20 stacks

Answer: A

Explanation:

AWS CloudFormation is an application management tool which provides application modelling, deployment, configuration, management and related activities. The limitations given below apply to the CloudFormation template and stack. There are no limits to the number of templates but each AWS CloudFormation account is limited to a maximum of 20 stacks by default. The Template, Parameter, Output, and Resource description fields are limited to 4096 characters. The user can include up to 60 parameters and 60 outputs in a template.

NEW QUESTION 181

- (Topic 3)

An organization has configured Auto Scaling for hosting their application. The system admin wants to understand the Auto Scaling health check process. If the instance is unhealthy, Auto Scaling launches an instance and terminates the unhealthy instance. What is the order execution?

- A. Auto Scaling launches a new instance first and then terminates the unhealthy instance
- B. Auto Scaling performs the launch and terminate processes in a random order
- C. Auto Scaling launches and terminates the instances simultaneously
- D. Auto Scaling terminates the instance first and then launches a new instance

Answer: D

Explanation:

Auto Scaling keeps checking the health of the instances at regular intervals and marks the instance for replacement when it is unhealthy. The ReplaceUnhealthy process terminates instances which are marked as unhealthy and subsequently creates new instances to replace them. This process first terminates the instance and then launches a new instance.

NEW QUESTION 185

- (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 188

- (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 checksu
- C. A success code is inserted into the S3 object metadat
- D. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successfu
- E. Amazon S3 is engineered for 99.999999999% durabilit
- F. Therefore there is no need to confirm that data was inserte

Answer: B

Explanation:

Reference:

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

NEW QUESTION 192

- (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 196

- (Topic 3)

A user has a weighing plant. The user measures the weight of some goods every 5 minutes and sends data to AWS CloudWatch for monitoring and tracking. Which of the below mentioned parameters is mandatory for the user to include in the request list?

- A. Value
- B. Namespace
- C. Metric Name
- D. Timezone

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 can publish the data to CloudWatch as single data points or as an aggregated set of data points called a statistic set. The user has to always include the namespace as part of the request. The user can supply a file instead of the metric name. If the user does not supply the timezone, it accepts the current time. If the user is sending the data as a single data point it will have parameters, such as value. However, if the user is sending as an aggregate it will have parameters, such as statistic-values.

NEW QUESTION 198

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 201

- (Topic 3)

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

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

Answer: AD

NEW QUESTION 203

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 208

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 211

- (Topic 3)

A sys admin is trying to understand the sticky session algorithm. Please select the correct sequence of steps, both when the cookie is present and when it is not, to help the admin understand the implementation of the sticky session:

ELB inserts the cookie in the response ELB chooses the instance based on the load balancing algorithm Check the cookie in the service request The cookie is found in the request The cookie is not found in the request

- A. 3,1,4,2 [Cookie is not Present] & 3,1,5,2 [Cookie is Present]
- B. 3,4,1,2 [Cookie is not Present] & 3,5,1,2 [Cookie is Present]
- C. 3,5,2,1 [Cookie is not Present] & 3,4,2,1 [Cookie is Present]
- D. 3,2,5,4 [Cookie is not Present] & 3,2,4,5 [Cookie is Present]

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. The load balancer uses a special load-balancer-generated cookie to track the application instance for each request. When the load balancer receives a request, it first checks to see if this cookie is present in the request. If so, the request is sent to the application instance specified in the cookie. If there is no cookie, the load balancer chooses an application instance based on the existing load balancing algorithm. A cookie is inserted into the response for binding subsequent requests from the same user to that application instance.

NEW QUESTION 216

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 217

- (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 221

- (Topic 3)

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

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

Answer: D

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. For ELB, CloudWatch provides various metrics including error code by ELB as well as by back-end servers (instances.. It gives data for the count of the number of HTTP response codes generated by the back-end instances. This metric does not include any response codes generated by the load balancer. These metrics are: The 2XX class status codes represents successful actions The 3XX class status code indicates that the user agent requires action The 4XX class status code represents client errors The 5XX class status code represents back-end server errors

NEW QUESTION 226

- (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 228

- (Topic 3)

A user has launched an EC2 instance. The instance got terminated as soon as it was launched. Which of the below mentioned options is not a possible reason for this?

- A. The user account has reached the maximum EC2 instance limit
- B. The snapshot is corrupt
- C. The AMI is missin
- D. It is the required part
- E. The user account has reached the maximum volume limit

Answer: A

Explanation:

When the user account has reached the maximum number of EC2 instances, it will not be allowed to launch an instance. AWS will throw an 'InstanceLimitExceeded' error. For all other reasons, such as "AMI is missing part", "Corrupt Snapshot" or "Volume limit has reached" it will launch an EC2 instance and then terminate it.

NEW QUESTION 231

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 233

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 238

- (Topic 3)

A user has configured ELB with two EBS backed instances. The user has stopped the instances for 1 week to save costs. The user restarts the instances after 1 week. Which of the below mentioned statements will help the user to understand the ELB and instance registration better?

- A. There is no way to register the stopped instances with ELB
- B. The user cannot stop the instances if they are registered with ELB
- C. If the instances have the same Elastic IP assigned after reboot they will be registered with ELB
- D. The instances will automatically get registered with ELB

Answer: C

Explanation:

Elastic Load Balancing registers the user's load balancer with his EC2 instance using the associated IP address. When the instances are stopped and started back they will have a different IP address. Thus, they will not get registered with ELB unless the user manually registers them. If the instances are assigned the same Elastic IP after reboot they will automatically get registered with ELB.

NEW QUESTION 240

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 243

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 245

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 247

- (Topic 3)

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

- A. Create an IAM policy with the security group and use that security group for AWS console login
- B. Create an IAM policy with a condition which denies access when the IP address range is not from the organization

- C. Configure the EC2 instance security group which allows traffic only from the organization's IP range
- D. Create an IAM policy with VPC and allow a secure gateway between the organization and AWS Console

Answer: B

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The user can add conditions as a part of the IAM policies. The condition can be set on AWS Tags, Time, and Client IP as well as on many other parameters. If the organization wants the user to access only from a specific IP range, they should set an IAM policy condition which denies access when the IP is not in a certain range. E.g. The sample policy given below denies all traffic when the IP is not in a certain range.

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

NEW QUESTION 252

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

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

Answer: A

Explanation:

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

NEW QUESTION 255

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

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

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