



Amazon-Web-Services

Exam Questions SOA-C01

AWS Certified SysOps Administrator - Associate

NEW QUESTION 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 2

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

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

Answer: B

NEW QUESTION 3

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

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

Answer: D

Explanation:

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

See also:

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

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

NEW QUESTION 4

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

Explanation:

By default: at no additional charge, Amazon RDS enables automated backups of your DB Instance with a 1-day retention period. By default: Amazon Redshift enables automated backups of your data warehouse cluster with a 1- day retention period.

NEW QUESTION 5

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: C

NEW QUESTION 6

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

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

- A. N
- B. Two instances in two different AZ's can't talk directly to each other via ICMP ping as that protocol is not allowed across subnet (i.e., broadcast) boundaries
- C. Ye
- D. Both the monitoring instance and the application instance have to be a part of the same security group, and that security group needs to allow inbound ICMP
- E. Ye
- F. The security group for the monitoring instance needs to allow outbound ICMP and the application instance's security group needs to allow Inbound ICMP
- G. Yes, Both the monitoring instance's security group and the application instance's security group need to allow both inbound and outbound ICMP ping packets since ICMP is not a connection- oriented protocol

Answer: C

NEW QUESTION 7

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

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

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

Answer: B

NEW QUESTION 8

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

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

Answer: A

Explanation:

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

NEW QUESTION 9

You are using ElastiCache Memcached to store session state and cache database queries in your infrastructure. You notice in CloudWatch that Evictions and GetMisses are both very high.
 What two actions could you take to rectify this? Choose 2 answers

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

Answer: AB

Explanation:

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

NEW QUESTION 10

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 10

Your organization's security policy requires that all privileged users either use frequently rotated passwords or one-time access credentials in addition to username/password.

Which two of the following options would allow an organization to enforce this policy for AWS users? Choose 2 answers

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

Answer: AB

Explanation:

See also: <http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html>

Enable MFA for privileged users

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

NEW QUESTION 15

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 of virtually unlimited size
- D. S3 allows you to store virtually unlimited amounts of data
- E. S3 offers Provisioned IOPS

Answer: AD

NEW QUESTION 17

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

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 20

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 24

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

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

- A. A High-Memory Quadruple Extra Large (m2.4xlarge) with EBS-Optimized set to true and a Provisioned IOPS EBS volume
- B. A Cluster Compute Eight Extra Large (cc2.8xlarge) using instance storage

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 28

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 (Mil)

Answer: C

Explanation:

When you launch an instance into a VPC, a primary private IP address from the address range of the subnet is assigned to the default network interface (eth0) of the instance. If you don't specify a primary private IP address, we select an available IP address in the subnet range for you

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

NEW QUESTION 33

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

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

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

Answer: BD

NEW QUESTION 36

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 37

A user has launched an EC2 instance. The user is planning to setup the CloudWatch alarm. Which of the below mentioned actions is not supported by the CloudWatch alarm?

- A. Notify the Auto Scaling launch config to scale up
- B. Send an SMS using SNS
- C. Notify the Auto Scaling group to scale down
- D. Stop the EC2 instance

Answer: A

Explanation:

A user can create a CloudWatch alarm that takes various actions when the alarm changes state. An alarm watches a single metric over the time period that the user has specified, and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The actions could be sending a notification to an Amazon Simple Notification Service topic (SMS, Email, and HTTP end point notifying the Auto Scaling policy or changing the state of the instance to Stop/Terminate.

CloudWatch cannot change the auto-scaling launch configuration.

B ?V It can send an SMS with SNS

C ?V Auto-scaling uses CloudWatch metrics to scale up and down.

D ?V CloudWatch can stop instances

NEW QUESTION 41

A user has setup a CloudWatch alarm on an EC2 action when the CPU utilization is above 75%. The alarm sends a notification to SNS on the alarm state. If the user wants to simulate the alarm action how can he achieve this?

- A. Run activities on the CPU such that its utilization reaches above 75%
- B. From the AWS console change the state to ??Alarm??
- C. The user can set the alarm state to ??Alarm?? using CLI
- D. Run the SNS action manually

Answer: C

Explanation:

Amazon CloudWatch alarms watch a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can test an alarm by setting it to any state using the SetAlarmState API (mon-set-alarm-state command). This temporary state change lasts only until the next alarm comparison occurs.
<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/AlarmThatSendsEmail.html>

NEW QUESTION 44

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 48

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

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

Answer: B

Explanation:

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

NEW QUESTION 49

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

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

Answer: A

Explanation:

Auto Scaling has its own aggregated CPU Utilization metric.

NEW QUESTION 54

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 58

A user has created an ELB with three instances. How many security groups will ELB create by default?

- A. 3
- B. 5
- C. 2
- D. 1

Answer: C

Explanation:

Elastic Load Balancing provides a special Amazon EC2 source security group that the user can use to ensure that back-end EC2 instances receive traffic only from Elastic Load Balancing. This feature needs two security groups: the source security group and a security group that defines the ingress rules for the back-end instances. To ensure that traffic only flows between the load balancer and the back-end instances, the user can add or modify a rule to the back-end security group which can limit the ingress traffic. Thus, it can come only from the source security group provided by Elastic Load Balancing.

NEW QUESTION 59

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

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

Answer: A

Explanation:

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

NEW QUESTION 63

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 68

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

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

Answer: C

Explanation:

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

NEW QUESTION 69

A user has configured the AWS CloudWatch alarm for estimated usage charges in the US East region. Which of the below mentioned statements is not true with respect to the estimated charges?

Exhibit:



- A. It will store the estimated charges data of the last 14 days
- B. It will include the estimated charges of every AWS service
- C. The metric data will represent the data of all the regions
- D. The metric data will show data specific to that region

Answer: D

Explanation:

When the user has enabled the monitoring of estimated charges for the AWS account with AWS CloudWatch, the estimated charges are calculated and sent several times daily to CloudWatch in the form of metric data. This data will be stored for 14 days. The billing metric data is stored in the US East (Northern Virginia) Region and represents worldwide charges. This data also includes the estimated charges for every service in AWS used by the user, as well as the estimated overall AWS charges.

NEW QUESTION 74

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 77

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 81

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 regio
- 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 83

A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created a public subnet CIDR (20.0.0.0/24. and VPN only subnets CIDR (20.0.1.0/24. along with the VPN gateway (vgw-12345. to connect to the user's data centre. Which of the below mentioned options is a valid entry for the main route table in this scenario?

- A. Destination: 20.0.0.0/24 and Target: vgw-12345
- B. Destination: 20.0.0.0/16 and Target: ALL
- C. Destination: 20.0.1.0/16 and Target: vgw-12345
- D. Destination: 0.0.0.0/0 and Target: vgw-12345

Answer: D

Explanation:

The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all traffic of the VPN subnet. Here are the valid entries for the main route table in this scenario: Destination: 0.0.0.0/0 & Target: vgw-12345 (To route all internet traffic to the VPN gateway.

Destination: 20.0.0.0/16 & Target: local (To allow local routing in VPC.

NEW QUESTION 88

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 92

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

- A. ELB sends data to CloudWatch every minute only and does not charge the user
- B. ELB will send data every minute and will charge the user extra
- C. ELB is not supported by CloudWatch
- D. It is not possible to setup detailed monitoring for ELB

Answer: A

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. Elastic Load Balancing includes 10 metrics and 2 dimensions, and sends data to CloudWatch every minute. This does not cost extra.

NEW QUESTION 95

A user has configured ELB with two EBS backed EC2 instances. The user is trying to understand the DNS access and IP support for ELB. Which of the below mentioned statements may not help the user understand the IP mechanism supported by ELB?

- A. The client can connect over IPV4 or IPV6 using Dualstack
- B. ELB DNS supports both IPV4 and IPV6
- C. Communication between the load balancer and back-end instances is always through IPV4
- D. The ELB supports either IPV4 or IPV6 but not both

Answer: D

Explanation:

Elastic Load Balancing supports both Internet Protocol version 6 (IPv6. and Internet Protocol version 4 (IPv4.. Clients can connect to the user's load balancer using either IPv4 or IPv6 (in EC2-Classic. DNS. However, communication between the load balancer and its back-end instances uses only IPv4. The user can use the Dualstack-prefixed DNS name to enable IPv6 support for communications between the client and the load balancers. Thus, the clients are able to access the load balancer using either IPv4 or IPv6 as their individual connectivity needs dictate.

NEW QUESTION 99

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 102

A user has configured Elastic Load Balancing by enabling a Secure Socket Layer (SSL) negotiation configuration known as a Security Policy. Which of the below mentioned options is not part of this secure policy while negotiating the SSL connection between the user and the client?

- A. SSL Protocols
- B. Client Order Preference
- C. SSL Ciphers
- D. Server Order Preference

Answer: B

Explanation:

Elastic Load Balancing uses a Secure Socket Layer (SSL) negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. A security policy is a combination of SSL Protocols, SSL Ciphers, and the Server Order Preference option.

NEW QUESTION 106

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 107

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 108

A user is planning to setup infrastructure on AWS for the Christmas sales. The user is planning to use Auto Scaling based on the schedule for proactive scaling. What advise would you give to the user?

- A. It is good to schedule now because if the user forgets later on it will not scale up
- B. The scaling should be setup only one week before Christmas
- C. Wait till end of November before scheduling the activity
- D. It is not advisable to use scheduled based scaling

Answer: C

Explanation:

Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. The user can specify any date in the future to

scale up or down during that period. As per Auto Scaling the user can schedule an action for up to a month in the future. Thus, it is recommended to wait until end of November before scheduling for Christmas.

NEW QUESTION 109

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 inflight requests continue to be served.

NEW QUESTION 112

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 115

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

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

Answer: D

Explanation:

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

NEW QUESTION 120

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 121

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

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

Answer: D

Explanation:

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

NEW QUESTION 123

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 127

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 131

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

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

Answer: C

Explanation:

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

NEW QUESTION 135

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

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

Answer: D

Explanation:

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

NEW QUESTION 136

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

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

Answer: C

Explanation:

Amazon RDS when enabled with Multi AZ will handle failovers automatically. Thus, the user can resume database operations as quickly as possible without administrative intervention. The primary DB instance switches over automatically to the standby replica if any of the following conditions occur:

- An Availability Zone outage
- The primary DB instance fails
- The DB instance's server type is changed
- The DB instance is undergoing software patching
- A manual failover of the DB instance was initiated using Reboot with failover

NEW QUESTION 139

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 144

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 145

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 150

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

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 153

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

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

Answer: A

Explanation:

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

NEW QUESTION 155

A user is observing the EC2 CPU utilization metric on CloudWatch. The user has observed some interesting patterns while filtering over the 1 week period for a particular hour. The user wants to zoom that data point to a more granular period. How can the user do that easily with CloudWatch?

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

Answer: A

NEW QUESTION 160

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 163

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

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

Answer: D

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. A user can have 5 IP addresses per region with EC2 Classic. The user can have 5 separate IPs with VPC in the same region as it has a separate limit than EC2 Classic.

NEW QUESTION 165

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 169

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

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

Answer: B

Explanation:

Amazon CloudWatch alarms watch a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The alarm has three states: Alarm, OK and Insufficient data. The Alarm will change to Insufficient Data when any of the three situations arise: when the alarm has just started, when the metric is not available or when enough data is not available for the metric to determine the alarm state. If the user wants to find that RDS is not available, he can setup to receive the notification when the state is in Insufficient data.

NEW QUESTION 170

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

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

Answer: A

Explanation:

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

NEW QUESTION 172

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 173

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

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

Answer: A

Explanation:

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

TLS 1.2
TLS 1.1
TLS 1.0
SSL 3.0
SSL 2.0

NEW QUESTION 177

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 180

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

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

Answer: B

Explanation:

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

NEW QUESTION 184

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 187

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 190

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 195

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 199

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

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

Answer: B

Explanation:

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

NEW QUESTION 204

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

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

Answer: B

Explanation:

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

NEW QUESTION 208

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 209

A storage admin wants to encrypt all the objects stored in S3 using server side encryption. The user does not want to use the AES 256 encryption key provided by S3. How can the user achieve this?

- A. The admin should upload his secret key to the AWS console and let S3 decrypt the objects
- B. The admin should use CLI or API to upload the encryption key to the S3 bucket
- C. When making a call to the S3 API mention the encryption key URL in each request
- D. S3 does not support client supplied encryption keys for server side encryption
- E. The admin should send the keys and encryption algorithm with each API call

Answer: D

Explanation:

AWS S3 supports client side or server side encryption to encrypt all data at rest. The server side encryption can either have the S3 supplied AES-256 encryption key or the user can send the key along with each API call to supply his own encryption key. Amazon S3 never stores the user's encryption key. The user has to supply it for each encryption or decryption call.

NEW QUESTION 211

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

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

Answer: D

Explanation:

AWS S3 provides a console to upload objects to a bucket. The user can use the file upload screen to upload the whole folder in one go by clicking on the Enable Enhanced Uploader option. When the user uploads 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 213

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 215

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 217

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 221

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 ?V 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 223

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

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

Answer: D

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. To enable detailed instance monitoring for a new Auto Scaling group, the user does not need to take any extra steps. When the user creates the AutoScaling launch config as the first step for creating an Auto Scaling group, each launch configuration contains a flag named InstanceMonitoring.Enabled. The default value of this flag is true. When the user has created a launch configuration with InstanceMonitoring.Enabled = false it will involve multiple steps to enable detail monitoring. The steps are:
Create a new Launch config with detailed monitoring enabled Update the Auto Scaling group with a new launch config Enable detail monitoring on each EC2 instance

NEW QUESTION 226

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

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

Answer: B

Explanation:

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

NEW QUESTION 231

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

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

Answer: D

Explanation:

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

NEW QUESTION 234

Your organization is preparing for a security assessment of your use of AWS.

In preparation for this assessment, which two IAM best practices should you consider implementing? Choose 2 answers

- A. Create individual IAM users for everyone in your organization
- B. Configure MFA on the root account and for privileged IAM users
- C. Assign IAM users and groups configured with policies granting least privilege access
- D. Ensure all users have been assigned and are frequently rotating a password, access ID/secret key, and X.509 certificate

Answer: BC

Explanation:

Reference:

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

NEW QUESTION 239

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

Answer: CD

NEW QUESTION 240

A company has an AWS account that contains three VPCs (Dev, Test, and Prod) in the same region.

Test is peered to both Prod and Dev. All VPCs have non-overlapping CIDR blocks. The company wants to push minor code releases from Dev to Prod to speed up time to market. Which of the following options helps the company accomplish this?

- A. Create a new peering connection Between Prod and Dev along with appropriate routes.
- B. Create a new entry to Prod in the Dev route table using the peering connection as the target.
- C. Attach a second gateway to De
- D. Add a new entry in the Prod route table identifying the gateway as the target.
- E. The VPCs have non-overlapping CIDR blocks in the same account
- F. The route tables contain local routes for all VPCs.

Answer: A

Explanation:

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/vpc-pg.pdf>

NEW QUESTION 241

An instance is launched into a VPC subnet with the network ACL configured to allow all inbound traffic and deny all outbound traffic. The instance's security group is configured to allow SSH from any IP address and deny all outbound traffic. What changes need to be made to allow SSH access to the instance?

- A. The outbound security group needs to be modified to allow outbound traffic.
- B. The outbound network ACL needs to be modified to allow outbound traffic.
- C. Nothing, it can be accessed from any IP address using SSH.
- D. Both the outbound security group and outbound network ACL need to be modified to allow outbound traffic.

Answer: B

Explanation:

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

NEW QUESTION 245

Which of the following are true regarding encrypted Amazon Elastic Block Store (EBS) volumes? Choose 2 answers

- A. Supported on all Amazon EBS volume types
- B. Snapshots are automatically encrypted
- C. Available to all instance types
- D. Existing volumes can be encrypted
- E. shared volumes can be encrypted

Answer: AB

Explanation:

This feature is supported on all Amazon EBS volume types (General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic). You can access encrypted Amazon EBS volumes the same way you access existing volumes; encryption and decryption are handled transparently and they require no additional action from you, your Amazon EC2 instance, or your application. Snapshots of encrypted Amazon EBS volumes are automatically encrypted, and volumes that are created from

encrypted Amazon EBS snapshots are also automatically encrypted.
Reference: <http://docs.aws.amazon.com/kms/latest/developerguide/services-ebs.html>

NEW QUESTION 246

A photo-sharing service stores pictures in Amazon Simple Storage Service (S3) and allows application sign-in using an OpenID Connect-compatible identity provider. Which AWS Security Token Service approach to temporary access should you use for the Amazon S3 operations?

- A. SAML-based Identity Federation
- B. Cross-Account Access
- C. AWS Identity and Access Management roles
- D. Web Identity Federation

Answer: D

NEW QUESTION 247

A customer is leveraging Amazon Simple Storage Service in eu-west-1 to store static content for a web-based property. The customer is storing objects using the Standard Storage class. Where are the customer's objects replicated?

- A. A single facility in eu-west-1 and a single facility in eu-central-1
- B. A single facility in eu-west-1 and a single facility in us-east-1
- C. Multiple facilities in eu-west-1
- D. A single facility in eu-west-1

Answer: C

NEW QUESTION 249

A company uses AWS Organization with a multi-account structure. A Syslog Administrator was notified that an IAM user with the System Administrator policy applied was not able to launch any Amazon EC2 instance using a public? Why is this occurring?

- A. The account is an AWS Organization master account, and by default it cannot provision EC2 instances.
- B. The account is an AWS Organization member account, and a service control policy is denying provisioning of EC2 instances.
- C. The account AWS Organization master account, and it does not have an access key activated for the IAM account.
- D. The account is an AWS Organization master account, and it does not have an access key activated for the IAM account.

Answer: B

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html

NEW QUESTION 250

A company must ensure that any objects upload to an bucket are encrypted. Which of the following actions will meet this requirement? (Select TWO.)

- A. Implement AWS Shield to protect against unencrypted objects stored in S3 buckets
- B. Implement Object access control list (ACL) to deny unencrypted objects from being uploaded to the S3 bucket.
- C. Implement Amazon S3 default encryption to make sure that any object being uploaded is encrypted before it is stored.
- D. Implement Amazon Inspector to inspect objects uploaded to the S3 I make sure that they are encrypted.
- E. Implement S3 bucket policies to deny unencrypted objects from being upload to the buckets.

Answer: BC

Explanation:

By default, all S3 buckets are private, and can only be accessed by users that have been explicitly granted access. Most use cases won't require broad-ranging public access to read files from your S3 buckets, unless you're using S3 to host public assets (for example, to host images for use on a public website), and it's best practice to never open access to the public. You can control access to your S3 resources by using a combination of bucket ACLs and IAM and bucket policies.

AWS also provides services that help you monitor and audit your security configurations, such as server access logging, Amazon CloudWatch Logs, AWS CloudTrail, and AWS Trusted Advisor.

NEW QUESTION 252

A SysOps Administrator must take a team's single existing AWS CloudFormation template and split it into smaller, service specific template. All of the service in the template reference a single, shared Amazon S3 bucket.

What should the Administrator do to ensure that this S3 bucket can be referenced by all the service templates?

- A. Include the S3 bucket as a mapping in each template
- B. Add the S3 bucket as a resource in each template
- C. Create the S3 bucket in its own template and export it
- D. Generate the S3 bucket using StackSets

Answer: D

NEW QUESTION 257

An organization created an Amazon Elastic File System (Amazon EFS) volume with a file system ID of fs-85baf1fc, and it is actively used by 10 Amazon EC2 hosts.

The organization has become concerned that the file system is not encrypted. How can this be resolved?

- A. Enable encryption on each hosts connection to the Amazon EFS volume Each connection must be recreated for encryption to take effect

- B. Enable encryption on the existing EFS volume by using the AWS Command Line Interface
- C. Enable encryption on each host's local drive Restart each host to encrypt the drive
- D. Enable encryption on a newly created volume and copy all data from the original volume Reconnect each host to the new volume

Answer: A

Explanation:

<https://docs.aws.amazon.com/efs/latest/ug/encryption.html> <https://aws.amazon.com/premiumsupport/knowledge-center/encrypt-data-efs/>

NEW QUESTION 260

An Amazon EC2 instance is unable to connect to an SMTP server in a different subnet. Other instances are successfully communication with the SMTP servers, however Flow Logs have been enabled on the SMTP server's network interface and show the following information

```
? 223342798652 eni-abc77deb 10.1.1.200 10.100.1.10 1123 25 17 70 48252 1515534437 1515535037 REJECT OK
```

- A. Add the instance to the security group for the SMTP server and ensure that it is permitted to communicate over TCP port 25.
- B. Disable the iptables server on the SMTP server so that the instance can properly communicate over the network.
- C. Install an email on the instance to ensure that it communicates correctly on TCP port 25 to theSMTP server.
- D. Add a rule to the security group for the instance to explicit permit TCP port 25 outbound to any address.

Answer: D

NEW QUESTION 264

After a particularly high bill, an organization wants to review the use of AWS services.

What AWS service will allow the SysOps Administrator to quickly view this information to shared it, and will also forest expenses for the billing period?

- A. AWS Trusted Advisor
- B. Amazon QuickSight
- C. AWS Cost and Usage Report
- D. AWS Cost Explorer

Answer: C

NEW QUESTION 267

A company application stores document within an Amazon S3 bucket. The application is running on Amazon EC3 in a VPC. A recent change in security requirement states traffic between the company's application and the S3 bucket must leave the Amazon network.

What AWS feature can provide this functionality?

- A. Security groups
- B. NAT gateways
- C. Virtual private gateway
- D. Gateway VPC endpoint

Answer: D

Explanation:

A VPC endpoint enables you to create a private connection between your VPC and another AWS service without requiring access over the Internet, through a NAT device, a VPN connection, or AWS Direct Connect. Endpoints are virtual devices.

NEW QUESTION 270

A SysOps Administrator needs a report of all IAM users and the status of MFA for each user. Which IAM feature would meet this requirement?

- A. IAM Rotes report
- B. IAM MFA report
- C. IAM User Policies report
- D. IAM Credential report

Answer: D

Explanation:

Getting Credential Reports for Your AWS Account

You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, and MFA devices. You can get a credential report from the AWS Management Console, the AWS SDKs and Command Line Tools, or the IAM API.

You can use credential reports to assist in your auditing and compliance efforts. You can use the report to audit the effects of credential lifecycle requirements, such as password and access key rotation. You can provide the report to an external auditor, or grant permissions to an auditor so that he or she can download the report directly.

You can generate a credential report as often as once every four hours. When you request a report, IAM first checks whether a report for the AWS account has been generated within the past four hours. If so, the most recent report is downloaded. If the most recent report for the account is older than four hours, or if there are no previous reports for the account, IAM generates and downloads a new report.

NEW QUESTION 273

A company has a new requirement stating that all resources in AWS must be tagged according to a set policy. Which AWS service should to enforce and continually identify all resources that are not in compliance with the policy?

- A. AWS CloudTrail
- B. Amazon Inspector
- C. AWS Config

D. AWS Systems Manager

Answer: C

Explanation:

AWS Config

TRACK RESOURCE INVENTORY AND CHANGES

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

NEW QUESTION 276

A SysOps Administrator must monitor a fleet of Amazon EC2 Linux instance with the constraint that no agent be installed. The SysOps administrator Chooses Amazon CloudWatch as the monitoring tool.

Which metrics can be measured given the constraints? (Select THREE.)

- A. CPU Utilization
- B. Disk Read Operations
- C. Memory Utilization
- D. Network Packets in
- E. Network Packets Dropped
- F. CPU Ready Time

Answer: ABD

Explanation:

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/viewing_metrics_with_cloudwatch.html

NEW QUESTION 277

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

Sequence	protocol	srcip	dstip	srcport	dstport	protocol	packets	bytes	start	end	action	log status
2	TCP	10.0.1.100	10.0.2.200	49786	3389	6	20	4349	1432917023	1432917042	ACCEPT	OK
2	TCP	10.0.2.200	10.0.1.100	3389	49786	6	20	5123	143291704	143291742	ACCEPT	OK

How can this issue be resolved?

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

Answer: D

NEW QUESTION 278

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 consol
- 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 280

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SOA-C01 Practice Exam Features:

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

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