



Amazon-Web-Services

Exam Questions SOA-C01

AWS Certified SysOps Administrator - Associate

NEW QUESTION 1

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

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 2

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

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

Answer: A

NEW QUESTION 3

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

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 4

The majority of your infrastructure is on-premises and you have a small footprint on AWS. Your company has decided to roll out a new application that is heavily dependent on low-latency connectivity to LDAP for authentication. Your security policy requires minimal changes to the company's existing application user management processes.

What option would you implement to successfully launch this application?

- A. Create a second, independent LDAP server in AWS for your application to use for authentication
- B. Establish a VPN connection so your applications can authenticate against your existing on-premises LDAP servers
- C. Establish a VPN connection between your data center and AWS, create a LDAP replica on AWS, and configure your application to use the LDAP replica for authentication
- D. Create a second LDAP domain on AWS, establish a VPN connection to establish a trust relationship between your new and existing domains, and use the new domain for authentication

Answer: C

Explanation:

Since it requires no changes to the authentication infrastructure as requested in the question, Option D creates a new LDAP, trusts, etc.

NEW QUESTION 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

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

Now you are not able to SSH into the host nor does it respond to web requests on port 80 that were working fine last time you checked. You have double-checked that all networking configuration parameters (security groups, route tables, IGW, EIP, NACLs, etc) are properly configured (and you haven't made any changes to those anyway since you were last able to reach the Instance). You look at the EC2 console and notice that system status check shows "impaired." Which should be your next step in troubleshooting and attempting to get the instance back to a healthy state so that you can log in again?

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

Answer: A

NEW QUESTION 7

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

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

Answer: AC

Explanation:

The below services provide Root level access:

- * EC2
- * Elastic Beanstalk
- * Elastic MapReduce ?V Master Node
- * Opswork

NEW QUESTION 8

You have a web-style application with a stateless but CPU and memory-intensive web tier running on a cc2 8xlarge EC2 instance inside of a VPC. The instance when under load is having problems returning requests within the SLA as defined by your business. The application maintains its state in a DynamoDB table, but the data tier is properly provisioned and responses are consistently fast. How can you best resolve the issue of the application responses not meeting your SLA?

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

Answer: C

Explanation:

But it is possibly A as DynamoDB is automatically available across three facilities in an AWS Region. So moving in to a same AZ is not possible / necessary. In this case the DB layer is not the issue, the EC2 8xlarge is the issue; so add another one with a ELB in front of it. See also: <https://aws.amazon.com/dynamodb/faqs/>

NEW QUESTION 9

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

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

Answer: C

NEW QUESTION 10

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

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

Answer: A

Explanation:

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

NEW QUESTION 10

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

Which deployment option meets these requirements while minimizing administrative burden?

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

Answer: A

Explanation:

OpsWorks has integrated support for Chef and lifecycle events.

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

NEW QUESTION 11

What are characteristics of Amazon S3? Choose 2 answers

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

Answer: AD

NEW QUESTION 13

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 15

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

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

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

Answer: B

Explanation:

Reference:

<https://aws.amazon.com/blogs/aws/amazon-s3-multipart-upload/>

NEW QUESTION 18

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 22

When an EC2 EBS-backed (EBS root) instance is stopped, what happens to the data on any ephemeral store volumes?

- A. Data will be deleted and will no longer be accessible
- B. Data is automatically saved in an EBS volume.
- C. Data is automatically saved as an EBS snapshot
- D. Data is unavailable until the instance is restarted

Answer: A

Explanation:

See: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html#instance-store-lifetime>

However, data in the instance store is lost under the following circumstances:

- ?V The underlying disk drive fails
- ?V The instance stops
- ?V The instance terminates

NEW QUESTION 26

Your company is moving towards tracking web page users with a small tracking image loaded on each page. Currently you are serving this image out of US-East, but are starting to get concerned about the time it takes to load the image for users on the west coast. What are the two best ways to speed up serving this image? Choose 2 answers.

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

Answer: AB

Explanation:

CloudFront gets the image closer to the user and Route53 ensures the best connection based on network latency. Option D does not address the issue.

NEW QUESTION 29

A user is trying to delete an Auto Scaling group from CLI. Which of the below mentioned steps are to be performed by the user?

- A. Terminate the instances with the `ec2-terminate-instance` command
- B. Terminate the Auto Scaling instances with the `as-terminate-instance` command
- C. Set the minimum size and desired capacity to 0
- D. There is no need to change the capacity
- E. Run the `as-delete-group` command and it will reset all values to 0

Answer: C

Explanation:

If the user wants to delete the Auto Scaling group, the user should manually set the values of the minimum and desired capacity to 0. Otherwise Auto Scaling will not allow for the deletion of the group from CLI. While trying from the AWS console, the user need not set the values to 0 as the Auto Scaling console will automatically do so.

NEW QUESTION 32

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

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

Answer: B

Explanation:

Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. In this case the load increases by Thursday and decreases by Friday. Thus, the user can setup the scaling activity based on the predictable traffic patterns of the web application using Auto Scaling scale by Schedule.

<http://docs.aws.amazon.com/cli/latest/reference/opsworks/set-time-based-auto-scaling.html>

NEW QUESTION 36

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 41

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

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

Answer: A

Explanation:

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

NEW QUESTION 42

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

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

Answer: B

Explanation:

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

NEW QUESTION 46

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

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

Answer: C

Explanation:

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

NEW QUESTION 48

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

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

Answer: D

Explanation:

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

NEW QUESTION 51

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

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

Answer: D

Explanation:

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

NEW QUESTION 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 a versa. The bucket policy can be defined at the bucket level which allows the objects as well as the bucket to be public with a single policy applied to that bucket. It cannot be applied at the object level. The folders are similar to objects with no content. Thus, folders can have only ACL and cannot have a policy.

NEW QUESTION 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 63

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

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

Answer: D

Explanation:

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

NEW QUESTION 68

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

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

Answer: B

Explanation:

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

NEW QUESTION 72

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 74

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 79

An organization is using AWS since a few months. The finance team wants to visualize the pattern of AWS spending. Which of the below AWS tool will help for this requirement?

- A. AWS Cost Manager
- B. AWS Cost Explorer
- C. AWS CloudWatch
- D. AWS Consolidated Billing

Answer: B

Explanation:

The AWS Billing and Cost Management console includes the Cost Explorer tool for viewing AWS cost data as a graph. It does not charge extra to user for this service. With Cost Explorer the user can filter graphs using resource tags or with services in AWS. If the organization is using Consolidated Billing it helps generate report based on linked accounts. This will help organization to identify areas that require further inquiry. The organization can view trends and use that to understand spend and to predict future costs.

NEW QUESTION 81

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

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

Answer: D

Explanation:

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

NEW QUESTION 86

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

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

Answer: A

Explanation:

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

NEW QUESTION 88

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 89

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

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

Answer: B

Explanation:

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

NEW QUESTION 94

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 95

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

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

Answer: B

Explanation:

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

NEW QUESTION 97

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

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

Answer: C

Explanation:

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy, eliminate I/O freezes, and minimize latency spikes during system backups. Running a DB instance with

high availability can enhance availability during planned system maintenance, and help protect your databases against DB instance failure and Availability Zone disruption. Note that the high-availability feature is not a scaling solution for read-only scenarios; you cannot use a standby replica to serve read traffic. To service read-only traffic, you should use a read replica.

NEW QUESTION 102

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

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

Answer: B

Explanation:

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

NEW QUESTION 103

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 106

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 108

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

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

Answer: D

Explanation:

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

NEW QUESTION 110

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

- A. GroupTotalInstances
- B. GroupSumInstances
- C. It is not possible to get a count of all the three metrics together

- D. The user has to find the individual number of running, terminating and pending instances and sum it
- E. GroupInstancesCount

Answer: A

Explanation:

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

NEW QUESTION 113

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

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

Answer: B

Explanation:

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

NEW QUESTION 115

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

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

Answer: A

Explanation:

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

NEW QUESTION 119

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 120

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

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

Answer: C

Explanation:

AWS provides cost allocation tags to categorize and track the AWS costs. When the user applies tags to his AWS resources, AWS generates a cost allocation report as a comma-separated value (CSV) file, with the usage and costs aggregated by those tags. Each tag will have a key-value and can be applied to services, such as EC2, S3, RDS, EMR, etc. The maximum size of a tag key is 128 Unicode characters.

NEW QUESTION 125

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

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

Answer: D

Explanation:

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

NEW QUESTION 127

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 132

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 135

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

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

Answer: C

Explanation:

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

NEW QUESTION 137

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

- A. Set up programmatic billing acces
- B. Download and parse the bill as per the requirement
- C. It is not possible for the user to create his own billing management service with AWS
- D. Enable the AWS CloudWatch alarm which will provide APIs to download the alarm data
- E. Use AWS billing APIs to download the usage report of each service from the AWS billing console

Answer: A

Explanation:

AWS provides an option to have programmatic access to billing. Programmatic Billing Access leverages the existing Amazon Simple Storage Service (Amazon S3) APIs. Thus, the user can build applications that reference his billing data from a CSV (comma-separated value) file stored in an Amazon S3 bucket. AWS will upload the bill to the bucket every few hours and the user can download the bill CSV from the bucket, parse it and create a billing system as per the requirement.

NEW QUESTION 140

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 144

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 146

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

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

Answer: D

Explanation:

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

NEW QUESTION 147

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

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

Answer: D

Explanation:

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

NEW QUESTION 150

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

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

Answer: C

Explanation:

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

NEW QUESTION 153

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

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

Answer: B

Explanation:

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

NEW QUESTION 157

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 162

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

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

Answer: C

Explanation:

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

NEW QUESTION 163

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 166

A sys admin has enabled logging on ELB. Which of the below mentioned fields will not be a part of the log file name?

- A. Load Balancer IP
- B. EC2 instance IP
- C. S3 bucket name
- D. Random string

Answer: B

Explanation:

Elastic Load Balancing access logs capture detailed information for all the requests made to the load balancer. Elastic Load Balancing publishes a log file from each load balancer node at the interval that the user has specified. The load balancer can deliver multiple logs for the same period. Elastic Load Balancing creates log file names in the following format:

??{Bucket}/{Prefix}/AWSLogs/{AWS AccountID}/elasticloadbalancing/{Region}/{Year}/{Month}/{Day}/{AWS Account ID}_elasticloadbalancing_{Region}_{Load Balancer Name}_{End Time}_{Load Balancer IP}_{Random String}.log??

NEW QUESTION 168

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

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

Answer: A

Explanation:

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

NEW QUESTION 172

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 177

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

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

Answer: A

Explanation:

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

NEW QUESTION 182

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 185

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

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

Answer: D

Explanation:

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

NEW QUESTION 186

A user has setup a CloudWatch alarm on the EC2 instance for CPU utilization. The user has setup to receive a notification on email when the CPU utilization is higher than 60%. The user is running a virus scan on the same instance at a particular time. The user wants to avoid receiving an email at this time. What should the user do?

- A. Remove the alarm
- B. Disable the alarm for a while using CLI
- C. Modify the CPU utilization by removing the email alert
- D. Disable the alarm for a while using the console

Answer: B

Explanation:

Amazon CloudWatch alarm watches a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. When the user has setup an alarm and it is known that for some unavoidable event the status may change to Alarm, the user can disable the alarm using the DisableAlarmActions API or from the command line `mon-disable-alarm-actions`.

NEW QUESTION 191

A root account owner is trying to understand the S3 bucket ACL. Which of the below mentioned options cannot be used to grant ACL on the object using the authorized predefined group?

- A. Authenticated user group
- B. All users group
- C. Log Delivery Group
- D. Canonical user group

Answer: D

Explanation:

An S3 bucket ACL grantee can be an AWS account or one of the predefined Amazon S3 groups. Amazon S3 has a set of predefined groups. When granting account access to a group, the user can specify one of the URLs of that group instead of a canonical user ID. AWS S3 has the following predefined groups: Authenticated Users group: It represents all AWS accounts. All Users group: Access permission to this group allows anyone to access the resource. Log Delivery group: WRITE permission on a bucket enables this group to write server access logs to the bucket.

NEW QUESTION 194

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

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

Answer: C

Explanation:

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

NEW QUESTION 196

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

- A. Stop the scaling process until research is completed
- B. It is not possible to find the root cause from that instance without triggering scaling
- C. Delete Auto Scaling until research is completed
- D. Suspend the scaling process until research is completed

Answer: D

Explanation:

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

NEW QUESTION 199

A user wants to find the particular error that occurred on a certain date in the AWS MySQL RDS DB. Which of the below mentioned activities may help the user to get the data easily?

- A. It is not possible to get the log files for MySQL RDS
- B. Find all the transaction logs and query on those records
- C. Direct the logs to the DB table and then query that table
- D. Download the log file to DynamoDB and search for the record

Answer: C

Explanation:

The user can view, download, and watch the database logs using the Amazon RDS console, the Command Line Interface (CLI) or the Amazon RDS API. For the MySQL RDS, the user can view the error log, slow query log, and general logs. The user can also view the MySQL logs easily by directing the logs to a database table in the main database and querying that table.

NEW QUESTION 204

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

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

Answer: C

Explanation:

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

NEW QUESTION 209

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

- A. AWS Auto Scaling

- B. AWS Route 53
- C. AWS EMR
- D. AWS SNS

Answer: B

Explanation:

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

NEW QUESTION 213

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

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

Answer: A

Explanation:

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

x-amz-server-side-encryption-customer-algorithm: Specifies the encryption algorithm

x-amz-server-side-encryption-customer-key: To provide the base64-encoded encryption key

x-amz-server-side-encryption-customer-key-MD5: To provide the base64-encoded 128-bit MD5 digest of the encryption key

NEW QUESTION 214

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

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

Answer: D

Explanation:

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

NEW QUESTION 217

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

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

Answer: D

Explanation:

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

NEW QUESTION 220

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 security policies is supported by ELB?

- A. Dynamic Security Policy
- B. All the other options
- C. Predefined Security Policy
- D. Default Security Policy

Answer: C

Explanation:

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

Predefined Security Policy, which comes with predefined cipher and SSL protocols; Custom Security Policy, which allows the user to configure a policy.

NEW QUESTION 221

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

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

Answer: D

Explanation:

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

NEW QUESTION 224

A user has created an application which will be hosted on EC2. The application makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK to connect with from the EC2 instance. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

- A. The user should attach an IAM role with DynamoDB access to the EC2 instance
- B. The user should create an IAM user with DynamoDB access and use its credentials within the application to connect with DynamoDB
- C. The user should create an IAM role, which has EC2 access so that it will allow deploying the application
- D. The user should create an IAM user with DynamoDB and EC2 access
- E. Attach the user with the application so that it does not use the root account credentials

Answer: A

Explanation:

With AWS IAM a user is creating an application which runs on an EC2 instance and makes requests to AWS, such as DynamoDB or S3 calls. Here it is recommended that the user should not create an IAM user and pass the user's credentials to the application or embed those credentials inside the application. Instead, the user should use roles for EC2 and give that role access to DynamoDB /S3. When the roles are attached to EC2, it will give temporary security credentials to the application hosted on that EC2, to connect with DynamoDB / S3.

NEW QUESTION 229

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 230

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

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

Answer: A

Explanation:

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

it will be single call to CloudWatch instead of 1000 calls if the data is aggregated.

NEW QUESTION 233

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

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

Answer: C

Explanation:

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

NEW QUESTION 237

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

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

Answer: D

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Whenever the organization is creating an IAM user, there should be a unique ID for each user. The names of users, groups, roles, instance profiles must be alphanumeric, including the following common characters: plus (+), equal (=), comma (,), period (.), at (@), and dash (-).

NEW QUESTION 242

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

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

Answer: C

Explanation:

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

NEW QUESTION 243

A user has created a VPC with public and private subnets using the VPC wizard. Which of the below mentioned statements is true in this scenario?

- A. The AWS VPC will automatically create a NAT instance with the micro size
- B. VPC bounds the main route table with a private subnet and a custom route table with a public subnet
- C. The user has to manually create a NAT instance
- D. VPC bounds the main route table with a public subnet and a custom route table with a private subnet

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. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance of a smaller or higher size, respectively. The VPC has an implied router and the VPC wizard updates the main route table used with the private subnet, creates a custom route table and associates it with the public subnet.

NEW QUESTION 247

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 252

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

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

Answer: B

Explanation:

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

NEW QUESTION 257

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 260

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 261

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 265

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

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

Answer: A

Explanation:

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

NEW QUESTION 267

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 271

How can you secure data at rest on an EBS volume?

- A. Encrypt the volume using the S3 server-side encryption service.
- B. Attach the volume to an instance using EC2's SSL interface.
- C. Create an IAM policy that restricts read and write access to the volume.
- D. Write the data randomly instead of sequentially.
- E. Use an encrypted file system on top of the EBS volume.

Answer: E

Explanation:

Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/policies_examples.html

NEW QUESTION 276

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

- A. Create a bucket policy and apply it to the bucket
- B. Create a NACL and attach it to the VPC of the bucket
- C. Create an ACL and apply it to all objects in the bucket
- D. Modify the IAM policies of any users that would access the bucket

Answer: A

Explanation:

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

NEW QUESTION 278

Your business is building a new application that will store its entire customer database on a RDS MySQL database, and will have various applications and users that will query that data for different purposes. Large analytics jobs on the database are likely to cause other applications to not be able to get the query results they need to, before time out. Also, as your data grows, these analytics jobs will start to take more time, increasing the negative effect on the other applications. How do you solve the contention issues between these different workloads on the same data?

- A. Enable Multi-AZ mode on the RDS instance
- B. Use ElastiCache to offload the analytics job data
- C. Create RDS Read-Replicas for the analytics work
- D. Run the RDS instance on the largest size possible

Answer: B

Explanation:

Amazon ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud. Amazon ElastiCache improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. The service simplifies and offloads the management, monitoring and operation of in-memory cache environments,

enabling your engineering resources to focus on developing applications. Using Amazon ElastiCache, you can not only improve load and response times to user actions and queries, but also reduce the cost associated with scaling web applications.

Amazon ElastiCache automates common administrative tasks required to operate a distributed cache environment. Using Amazon ElastiCache, you can add a caching layer to your application architecture in a matter of minutes via a few clicks of the AWS Management Console. Once a cache cluster is provisioned, Amazon ElastiCache automatically detects and replaces failed cache nodes, providing a resilient system that mitigates the risk of overloaded databases, which slow website and application load times. Through integration with Amazon CloudWatch monitoring, Amazon ElastiCache provides enhanced visibility into key performance metrics associated with your cache nodes. Amazon ElastiCache is protocol-compliant with Memcached and Redis, so code, applications, and popular tools that you use today with your existing Memcached or Redis environments will work seamlessly with the service. As with all Amazon Web Services,

NEW QUESTION 283

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 286

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 287

A customer needs to capture all client connection information from their load balancer every five minutes. The company wants to use this data for analyzing traffic patterns and troubleshooting their applications. Which of the following options meets the customer requirements?

- A. Enable AWS CloudTrail for the load balancer.
- B. Enable access logs on the load balancer.
- C. Install the Amazon CloudWatch Logs agent on the load balancer.
- D. Enable Amazon CloudWatch metrics on the load balancer.

Answer: A

NEW QUESTION 292

Which features can be used to restrict access to data in S3? Choose 2 answers

- A. Set an S3 ACL on the bucket or the object.
- B. Create a CloudFront distribution for the bucket.
- C. Set an S3 bucket policy.
- D. Enable IAM Identity Federation
- E. Use S3 Virtual Hosting

Answer: AC

Explanation:

<https://aws.amazon.com/s3/faqs/>

NEW QUESTION 296

The Database Administrator learn is interested in performing manual backups of Amazon DRS Oracle DB instance. What step be taken to perform the backups?

- A. Attach an Amazon EBS volume with Oracle RMAN installed to the RDS instance
- B. Take a snapshot of the EBS volume that is attached to the DB instance.
- C. Install Oracle Secure Backup on the RDS instance and back up the Oracle database to Amazon S3
- D. Take a snapshot of the DB instance

Answer: D

NEW QUESTION 301

A SysOps Administrator needs to implement logging strategy that will allow of Linux-based Amazon EC2 instance to write log files into a single shared archive. An

additional requirement is that log location must be accessible on all EC2 fleet instances using the local file system. What service meets the requirements?

- A. Amazon Elastic IV
- B. Amazon EBS
- C. Amazon Kinesis
- D. AWS CloudTrail

Answer: C

Explanation:

Amazon Kinesis Firehose

Customers who have large amounts of log data to process can use Amazon Kinesis Firehose as a serverless log ingestion and delivery mechanism. Amazon Kinesis Firehose is a managed service that enables customers to deliver real-time streaming data to destinations such as Amazon ES, Amazon S3, and Amazon Redshift. Firehose is designed to handle large amounts of incoming data and can generate bulk indexing requests to an Amazon ES domain.

Unlike self-managed log processing components, such as a Logstash cluster, Firehose does not require any servers, applications, or resource management. Customers configure individual data producers to send log data to a Firehose delivery stream continuously, and Firehose manages the rest.

NEW QUESTION 305

A new network is needed to run secure Amazon EC2 instance. This network cannot have direct access to the internet and must be separate from existing production instances. The instances will be manager using SSH from a Developer in a home office with a fixed IP address but without a VPN- capable router. How should a SysOps Administrator create this network and manage these servers?

- A. Create a new subnet in an existing VP
- B. Configure access rules to allow SSH access from the Developer's IP address
- C. Use AWS Shield to select the instances that should not have access to the internet.
- D. Associated an internet gateway with a new VPC with two subne
- E. Set up a bastion instance with an Elastic IP address Configure security groups and routing to allow SSH access to the bastion instance from the Developer's Ip address and SSH access from the bastion hot to the private subnet.
- F. Configure a new VPC with one public subnet no internet gatewa
- G. Configure the security for the instance to allow SSH from the Developer's IP address.
- H. Setup a new VPC with one private subne
- I. When deployment the instance use the User data to install and configure a third-party management tool for the instances Connect to the instance using the third-party tool.

Answer: C

NEW QUESTION 310

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 312

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

```
2 223342758652 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 313

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 316

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 320

A SysOps Administrator has set up a new Application Load Balancer (ALB) in front of a pair of private web server in multiple Availability Zones. After deployment an updates CloudFormation template with many changes, user now goes to one web server only. What is the MOST likely reason that the traffic is not being balanced between both servers?

- A. The faulty is returning HTTP 200 has been removed.
- B. Sticky session have been disabled in the ALB for the working sever.
- C. The ALB using a custom ping path that is not found on the faulty server.
- D. The web client are using HTTP/2, which is terminated at the ALB.

Answer: B

Explanation:

Until now, the behavior of load balancers has been to route each request independently to the Amazon EC2 instance with the least load. With the stickiness feature, you can configure the load balancer to bind user sessions to specific application instances. All requests coming from the user during the session will be sent to the same application instance. Elastic Load Balancing supports two mechanisms to provide session stickiness: load balancer-generated HTTP cookies, which allow browser-based session lifetimes, and application-generated HTTP cookies, which allow application-specific session lifetimes. You can learn more about this feature by visiting the ELB Developers Guide.

NEW QUESTION 325

A SysOps Administrator management a fleet of instance store-backed Amazon Linux EC2 instances. The SSH key used to access these instances has been lost. How can SSH access be restored?

- A. Contact AWS Support to retrieve a backup of the SSH key after authentication
- B. Create a new SSH key stop the EC2 instances apply the new key, and restart the EC2 instances
- C. Create a new SSH key and apply the new key to the running EC2 instances
- D. Launch a new fleet of EC2 instances with a newly created SSH key

Answer: A

Explanation:

Resolution

Warning: Do not perform this procedure if your EC2 instance is an instance store-backed instance. This recovery procedure requires a stop and start of your instance, which means that data on instance store volumes will be lost. For more information, see Determining the Root Device Type of Your Instance. To recover access to your Linux instance using AWS Systems Manager (SSM) automation, run the AWSSupport-ResetAccess Automation automation document. For more information, see Reset Passwords and SSH Keys on Amazon EC2 Instances. Or, to manually recover access to your Linux instance, create a new key pair to replace the lost key pair. For more information, see Connecting to Your Linux Instance If You Lose Your Private Key.

NEW QUESTION 326

An organization must reduce unexpected overages in spending, and a SysOps administrator has been asked to automate a report of the AWS environment. How can this be accomplished with MINIMAL effort?

- A. Set spending limits in AWS Trusted Advisor and enable weekly email reports.
- B. Set up an AWS Lambda function to send AWS Cost and Usage reports through email
- C. Use Amazon Inspector to alert if predefined spending threshold are exceeded.
- D. Configure a customer budget with email alert through API actions for AWS Budgets.

Answer: D

Explanation:

<https://docs.aws.amazon.com/sdk-for-go/api/service/budgets/>

NEW QUESTION 327

A System Administrator is trying to identify why Put Object calls are not made from an Amazon EC2 instance to an Amazon bucket in the same region. The instance is launched in a subnet with CIDR range 10.1.0.24 and 'Auto assign public IP set to yes. The instance profile tied to this instance has AmazonS3Access policy. Security group rules for the instance:

```

Protocol      PortRange    Source
HTTP (80)    80           0.0.0.0/0
HTTPS (443)  443          0.0.0.0/0
Custom TCP   1024-65535  0.0.0.0/0

The route table for the subnet in which this instance is launched
Destination  Target
10.0.0.0/16  local
    
```

Based on the information provided what is causing the lack of access to S3 from the instance?

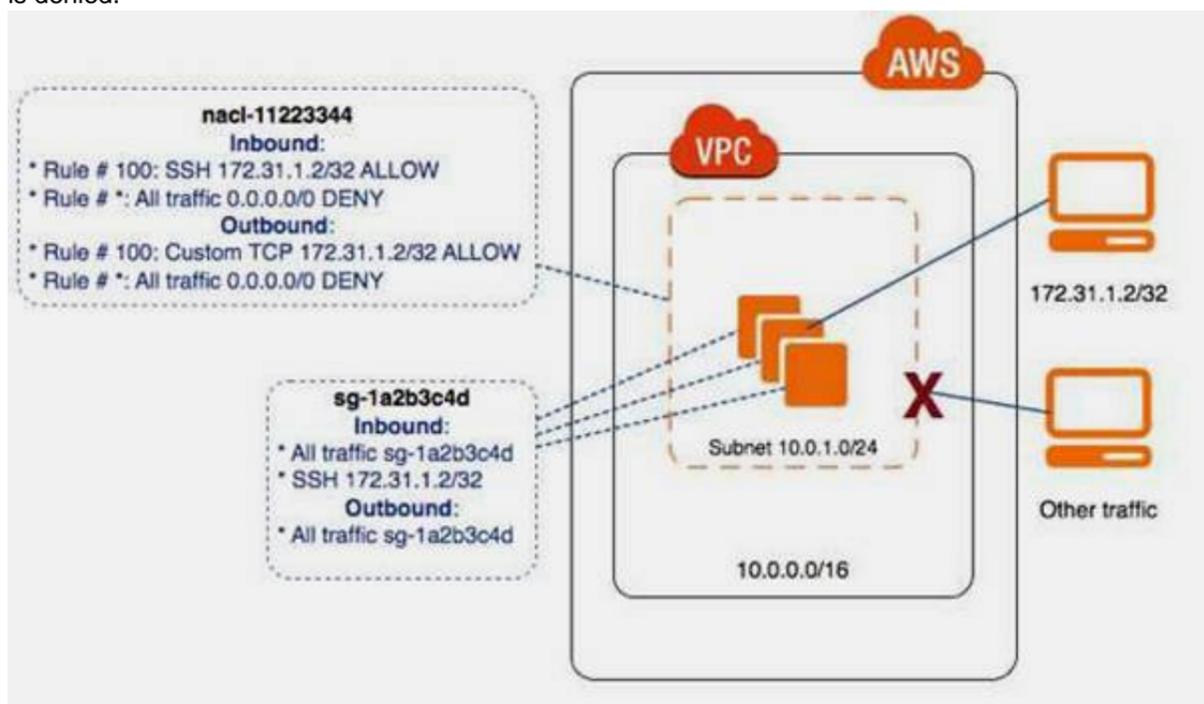
- A. The instance profile does not have explicit permissions to write objects to the S3 bucket.
- B. The route table does not have a rule for all traffic to pass through a NAT gateway.
- C. The route table does not have rule for all traffic to pass through an internet gateway

Answer: B

Explanation:

Controlling Access to Instances in a Subnet

In this example, instances in your subnet can communicate with each other, and are accessible from a trusted remote computer. The remote computer may be a computer in your local network or an instance in a different subnet or VPC that you use to connect to your instances to perform administrative tasks. Your security group rules and network ACL rules allow access from the IP address of your remote computer (172.31.1.2/32). All other traffic from the Internet or other networks is denied.



All instances use the same security group (sg-1a2b3c4d), with the following rules.

Protocol Protocol Port Source Comments

Type	Range			
All traffic	All	All	sg-1a2b3c4d	Enables instances associated with the same security group to communicate with each other.
TCP	SSH	22	172.31.1.2/32	Allows inbound SSH access from the remote computer. If the instance is a Windows computer, then this rule must use the RDP protocol for port 3389 instead.

Protocol Type	Protocol	Port Range	Destination	Comments
All traffic	All	All	sg-1a2b3c4d	Enables instances associated with the same security group to communicate with each other.

The subnet is associated with a network ACL that has the following rules.

Rule #	Type	Protocol	Port Range	Source	Allow/Deny	Comments
100	SSH	TCP	22	172.31.1.2/32	ALLOW	Allows inbound traffic from the remote computer. If the instance is a Windows computer, then this rule must use the RDP protocol for port 3389 instead.
*	All traffic	All	All	0.0.0.0/0	DENY	Denies all other inbound traffic that does not match the previous rule.

Rule #	Type	Protocol	Port Range	Destination	Allow/Deny	Comments
100	Custom TCP	TCP	1024-65535	172.31.1.2/32	ALLOW	Allows outbound responses to the remote computer. Network ACLs are stateless, therefore this rule is required to allow response traffic for inbound requests.
*	All traffic	All	All	0.0.0.0/0	DENY	Denies all other outbound traffic that does not match the previous rule.

This scenario gives you the flexibility to change the security groups or security group rules for your instances, and have the network ACL as the backup layer of defense. The network ACL rules apply to all instances in the subnet, so if you accidentally make your security group rules too permissive, the network ACL rules

continue to permit access only from the single IP address. For example, the following rules are more permissive than the earlier rules ?X they allow inbound SSH access from any IP address.

Type	Protocol	Port Range	Source	Comments
All traffic	All	All	sg-1a2b3c4d	Enables instances associated with the same security group to communicate with each other.
SSH	TCP	22	0.0.0.0/0	Allows SSH access from any IP address.

Type	Protocol	Port Range	Destination	Comments
All traffic	All	All	0.0.0.0/0	Allows all outbound traffic.

However, only other instances within the subnet and your remote computer are able to access this instance. The network ACL rules still prevent all inbound traffic to the subnet except from your remote computer.

NEW QUESTION 328

A company has a VoIP application deployed on AWS. The application is accessed by employees in a remote office and is extremely sensitive to any latency and packets loss. Minimize latency and packet loss is a higher priority than minimizing cost.

Employees are reporting occasional difficulties accessing the application. The Local Network Engineer has completed thorough troubleshooting on the LAN and unable to identify any signs of congestion or equipment failure that may be causing the issue.

What is the BEST way to address the connectivity issues between the remote office and the application?

- A. Configure a VPN connection to the VPC Route all traffic to the application via the VPN connection over the public internet
- B. Establish a Direct Connect to the VPC Route all traffic to the application via the direct connect connection
- C. Enable VPC peering to decrease latency between instances Enable QoS on peering connection
- D. Configure Amazon Trusted Advisor to give higher prioritization to the IP to assigned to the remote office over public internet traffic

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpc/latest/peering/create-vpc-peering-connection.html>

NEW QUESTION 331

.....

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](#)