

# Amazon-Web-Services

## Exam Questions SOA-C01

AWS Certified SysOps Administrator - Associate



#### NEW QUESTION 1

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

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

- A. Create an AD policy to modify Windows Firewall settings on all hosts in the VPC to deny access from the IP address block
- B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP address block
- C. Add a rule to all of the VPC 5 Security Groups to deny access from the IP address block
- D. Modify the Windows Firewall settings on all Amazon Machine Images (AMIs) that your organization uses in that VPC to deny access from the IP address block

**Answer:** B

#### Explanation:

Reference:

[http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_SecurityGroups.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html)

#### NEW QUESTION 2

Which two AWS services provide out-of-the-box user configurable automatic backup-as-a-service and backup rotation options? Choose 2 answers

- A. Amazon S3
- B. Amazon RDS
- C. Amazon EBS
- D. Amazon Redshift

**Answer:** BD

#### Explanation:

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

#### NEW QUESTION 3

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 users credentials into the instance User Data
- C. Modify the appropriate Cloud Watch 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 4

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 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

What is a placement group?

- A. A collection of Auto Scaling groups in the same Region
- B. Feature that enables EC2 instances to interact with each other via high bandwidth, low latency connections
- C. A collection of Elastic Load Balancers in the same Region or Availability Zone

D. A collection of authorized Cloud Front edge locations for a distribution

**Answer: B**

**Explanation:**

Reference:

<http://aws.amazon.com/ec2/faqs/>

A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gigabits per second (Gbps) network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both

**NEW QUESTION 7**

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

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

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

**Answer: C**

**NEW QUESTION 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**

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

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

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

**Answer: B**

**NEW QUESTION 10**

You are running a web-application on AWS consisting of the following components an Elastic Load Balancer (ELB) an Auto-Scaling Group of EC2 instances running Linux/PHP/Apache, and Relational Database Service (RDS) MySQL.

Which security measures fall into AWS's responsibility?

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

**Answer: B**

#### NEW QUESTION 10

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

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

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

**Answer: C**

#### Explanation:

Reference:

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

#### NEW QUESTION 15

Your team is excited about the use of AWS because now they have access to "programmable Infrastructure". You have been asked to manage your AWS infrastructure in a manner similar to the way you might manage application code. You want to be able to deploy exact copies of different versions of your infrastructure, stage changes into different environments, revert back to previous versions, and identify what versions are running at any particular time (development, test, QA, production).

Which approach addresses this requirement?

- A. Use cost allocation reports and AWS OpsWorks to deploy and manage your infrastructure.
- B. Use AWS CloudWatch metrics and alerts along with resource tagging to deploy and manage your infrastructure.
- C. Use AWS Beanstalk and a version control system like GIT to deploy and manage your infrastructure.
- D. Use AWS CloudFormation and a version control system like GIT to deploy and manage your infrastructure.

**Answer: D**

#### Explanation:

Reference:

?V Answer A: does not provide versioning

?V Answer B: does not provide versioning

?V Answer C: Beanstalk provide version control over your application (not infrastructure)

Extract from what is AWS CloudFormation: (<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html>)

Easily Control and Track Changes to Your Infrastructure In some cases, you might have underlying resources that you want to upgrade incrementally. For example, you might change to a higher performing instance type in your Auto Scaling launch configuration so that you can reduce the maximum number of instances in your Auto Scaling group. If problems occur after you complete the update, you might need to roll back your infrastructure to the original settings. To do this manually, you not only have to remember which resources were changed, you also have to know what the original settings were.

When you provision your infrastructure with AWS CloudFormation, the AWS CloudFormation template describes exactly what resources are provisioned and their settings. Because these templates are text files, you simply track differences in your templates to track changes to your infrastructure, similar to the way developers control revisions to source code. For example, you can use a version control system with your templates so that you know exactly what changes were made, who made them, and when. If at any point you need to reverse changes to your infrastructure, you can use a previous version of your template.

#### NEW QUESTION 18

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

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 21

An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

- A. The organization should create a separate login ID but give the IAM users the same alias so that each one can login with their alias
- B. The organization should create each user in a separate region so that they have their own URL to login
- C. It is not possible to have the same login ID for multiple IAM users of the same account
- D. The organization should create various groups and add each user with the same login ID to different group
- E. The user can login with their own group ID

**Answer: C**

#### Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Whenever the organization is creating an IAM user, there should be a unique ID for each user. It is not possible to have the same login ID for multiple users. The names of users, groups, roles, instance profiles must be alphanumeric, including the following common characters: plus (+), equal (=), comma (,), period (.), at (@), and dash (-).

#### NEW QUESTION 23

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

- A. AWS Simple Notification Service
- B. AWS Simple Workflow
- C. AWS Simple Queue Service
- D. AWS Simple Query Service

**Answer:** C

**Explanation:**

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

**NEW QUESTION 26**

An organization has created 50 IAM users. The organization has introduced a new policy which will change the access of an IAM user. How can the organization implement this effectively so that there is no need to apply the policy at the individual user level?

- A. Use the IAM groups and add users as per their role to different groups and apply policy to group
- B. The user can create a policy and apply it to multiple users in a single go with the AWS CLI
- C. Add each user to the IAM role as per their organization role to achieve effective policy setup
- D. Use the IAM role and implement access at the role level

**Answer:** A

**Explanation:**

With AWS IAM, a group is a collection of IAM users. A group allows the user to specify permissions for a collection of users, which can make it easier to manage the permissions for those users. A group helps an organization manage access in a better way; instead of applying at the individual level, the organization can apply at the group level which is applicable to all the users who are a part of that group.

**NEW QUESTION 27**

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 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 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 37

A sysadmin has created the below mentioned policy and applied to an S3 object named aws.jpg. The aws.jpg is inside a bucket named cloudacademy. What does this policy define?

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

- A. It is not possible to define a policy at the object level
- B. It will make all the objects of the bucket cloudacademy as public
- C. It will make the bucket cloudacademy as public
- D. the aws.jpg object as public

**Answer:** A

#### NEW QUESTION 40

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 43

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 44

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

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

**Answer:** A

#### Explanation:

Auto Scaling has its own aggregated CPU Utilization metric.

#### NEW QUESTION 47

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

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

**Answer:** B

#### Explanation:

[http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported\\_services.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html) CloudWatch monitors the following services. As soon as you begin using a service, it automatically sends metrics to CloudWatch for you. CloudWatch offers either basic or detailed monitoring for supported AWS products. Basic monitoring means that a service sends data points to CloudWatch every five minutes. Detailed monitoring means that a service sends data points to CloudWatch every minute.

**Note**

If you are using a service that supports both basic and detailed data collection (for example, Amazon EC2 and Auto Scaling), and you want to access detailed statistics, you must enable detailed metric collection for that service.

**Auto Scaling**

Auto Scaling sends data to CloudWatch every 5 minutes by default. For an additional charge, you can enable detailed monitoring for Auto Scaling, which sends data to CloudWatch every minute. You can create alarms using Auto Scaling Dimensions and Metrics. For more information, see Monitor Your Auto Scaling Instances in the Auto Scaling User Guide.

**NEW QUESTION 48**

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

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

**Answer: B**

**Explanation:**

AWS EBS supports encryption of the volume while creating new volumes. It also supports creating volumes from existing snapshots provided the snapshots are created from encrypted volumes. The data at rest, the I/O as well as all the snapshots of EBS will be encrypted. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between the EC2 instances and EBS storage. EBS encryption is based on the AES-256 cryptographic algorithm, which is the industry standard.

**NEW QUESTION 53**

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 54**

A user has a refrigerator plant. The user is measuring the temperature of the plant every 15 minutes. If the user wants to send the data to CloudWatch to view the data visually, which of the below mentioned statements is true with respect to the information given above?

- A. The user needs to use AWS CLI or API to upload the data
- B. The user can use the AWS Import Export facility to import data to CloudWatch
- C. The user will upload data from the AWS console
- D. The user cannot upload data to CloudWatch since it is not an AWS service metric

**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. While sending the data the user has to include the metric name, namespace and timezone as part of the request.

**NEW QUESTION 56**

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 60**

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

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

**Answer: C**

**Explanation:**

A system admin can grant permission of the S3 objects or buckets to any user or make the objects public using the bucket policy and user policy. Both use the JSON-based access policy language.

Generally, if the user is defining the ACL on the bucket, the objects in the bucket do not inherit it and vice versa. The bucket policy can be defined at the bucket level which allows the objects as well as the bucket to be public with a single policy applied to that bucket.

**NEW QUESTION 61**

A root AWS account owner is trying to understand various options to set the permission to AWS S3. Which of the below mentioned options is not the right option to grant permission for S3?

- A. User Access Policy
- B. S3 Object Access Policy
- C. S3 Bucket Access Policy
- D. S3 ACL

**Answer: B**

**Explanation:**

Amazon S3 provides a set of operations to work with the Amazon S3 resources. Managing S3 resource access refers to granting others permissions to work with S3. There are three ways the root account owner can define access with S3:

S3 ACL: The user can use ACLs to grant basic read/write permissions to other AWS accounts.

S3 Bucket Policy: The policy is used to grant other AWS accounts or IAM users permissions for the bucket and the objects in it.

User Access Policy: Define an IAM user and assign him the IAM policy which grants him access to S3.

**NEW QUESTION 66**

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 71**

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

Exhibit:



- A. It will store the estimated charges data of the last 14 days

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 74**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 79**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 84**

A user has created a queue named myqueue with SQS. There are four messages published to queue which are not received by the consumer yet. If the user tries to delete the queue, what will happen?

- A. A user can never delete a queue manually
- B. AWS deletes it after 30 days of inactivity on queue
- C. It will delete the queue
- D. It will initiate the delete but wait for four days before deleting until all messages are deleted automatically.
- E. It will ask user to delete the messages first

**Answer:** B

**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. The user can delete a queue at any time, whether it is empty or not. It is important to note that queues retain messages for a set period of time. By default, a queue retains messages for four days.

**NEW QUESTION 86**

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 90

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

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 93

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

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

**Answer: A**

#### Explanation:

CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. Elastic Load Balancing includes 10 metrics and 2 dimensions, and sends data to CloudWatch every minute. This does not cost extra.

#### NEW QUESTION 95

A user is checking the CloudWatch metrics from the AWS console. The user notices that the CloudWatch data is coming in UTC. The user wants to convert the data to a local time zone. How can the user perform this?

- A. In the CloudWatch dashboard the user should set the local timezone so that CloudWatch shows the data only in the local time zone
- B. In the CloudWatch console select the local timezone under the Time Range tab to view the data as per the local timezone
- C. The CloudWatch data is always in UTC; the user has to manually convert the data
- D. The user should have send the local timezone while uploading the data so that CloudWatch will show the data only in the local timezone

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 96

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

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

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

**Answer: D**

#### Explanation:

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

```
"Statement": [
{
  "Sid": "AllowUsersAllActionsForCredentials", "Effect": "Allow",
  "Action": [ "iam:*AccessKey*",
  ],

```

```
"Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
}
]
```

#### NEW QUESTION 100

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 101

A user has created a queue named ??myqueue?? in US-East region with AWS SQS. The user??s AWS account ID is 123456789012. If the user wants to perform some action on this queue, which of the below Queue URL should he use?

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 106

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 108

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 113

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

- A. AWS CloudWatch + AWS SES
- B. AWS CloudWatch + AWS SNS

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

**Answer:** B

**Explanation:**

Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and fully managed push messaging service. Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. The user can configure some sensor devices at his home which receives data on the HTTP end point (REST calls) and turn on the red light. The user can configure the CloudWatch alarm to send a notification to the AWS SNS HTTP end point (the sensor device) and it will turn the light red when there is an alarm condition.

**NEW QUESTION 116**

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 117**

A user is running one instance for only 3 hours every day. The user wants to save some cost with the instance. Which of the below mentioned Reserved Instance categories is advised in this case?

- A. The user should not use RI; instead only go with the on-demand pricing
- B. The user should use the AWS high utilized RI
- C. The user should use the AWS medium utilized RI
- D. The user should use the AWS low utilized RI

**Answer:** A

**Explanation:**

The AWS Reserved Instance provides the user with an option to save some money by paying a one-time fixed amount and then save on the hourly rate. It is advisable that if the user is having 30% or more usage of an instance per day, he should go for a RI. If the user is going to use an EC2 instance for more than 2200-2500 hours per year, RI will help the user save some cost. Here, the instance is not going to run for less than 1500 hours. Thus, it is advisable that the user should use the on-demand pricing.

**NEW QUESTION 122**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 127**

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

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

**Answer:** D

**Explanation:**

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

the organization network into the temporary AWS security credentials. This proxy server will attach an IAM role to the user after authentication.

#### NEW QUESTION 130

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

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

**Answer: C**

#### Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. Security groups work at the instance level. If the user is using the default security group it will have a rule which allows the instances to communicate with other. For a new security group the user has to specify the rule, add it to define the source as the security group itself, and select all the protocols and ports for that source.

#### NEW QUESTION 131

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

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

**Answer: A**

#### Explanation:

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

#### NEW QUESTION 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 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 140

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

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

D. It will wait for the user's input about the error and correct the mistake after the input

**Answer:** C

**Explanation:**

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

**NEW QUESTION 141**

An organization is measuring the latency of an application every minute and storing data inside a file in the JSON format. The organization wants to send all latency data to AWS CloudWatch. How can the organization achieve this?

- A. The user has to parse the file before uploading data to CloudWatch
- B. It is not possible to upload the custom data to CloudWatch
- C. The user can supply the file as an input to the CloudWatch command
- D. The user can use the CloudWatch Import command to import data from the file to CloudWatch

**Answer:** C

**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 part of the request. If the user wants to upload the custom data from a file, he can supply file name along with the parameter -- metric-data to command put-metric-data.

**NEW QUESTION 143**

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 148**

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

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

**Answer:** C

**Explanation:**

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

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

**NEW QUESTION 150**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 151**

A user has enabled session stickiness with ELB. The user does not want ELB to manage the cookie; instead he wants the application to manage the cookie. What will happen when the server instance, which is bound to a cookie, crashes?

- A. The response will have a cookie but stickiness will be deleted
- B. The session will not be sticky until a new cookie is inserted
- C. ELB will throw an error due to cookie unavailability
- D. The session will be sticky and ELB will route requests to another server as ELB keeps replicating the Cookie

**Answer: B**

**Explanation:**

With Elastic Load Balancer, if the admin has enabled a sticky session with application controlled stickiness, the load balancer uses a special cookie generated by the application to associate the session with the original server which handles the request. ELB follows the lifetime of the application-generated cookie corresponding to the cookie name specified in the ELB policy configuration. The load balancer only inserts a new stickiness cookie if the application response includes a new application cookie. The load balancer stickiness cookie does not update with each request. If the application cookie is explicitly removed or expires, the session stops being sticky until a new application cookie is issued.

**NEW QUESTION 152**

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

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

**Answer: A**

**NEW QUESTION 154**

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 155**

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

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

**Answer: D**

**Explanation:**

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

**NEW QUESTION 159**

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 161

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

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

**Answer:** A

#### Explanation:

Elastic Load Balancing uses a Secure Socket Layer (SSL) negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. If the load balancer is configured to support the Server Order Preference, then the load balancer gets to select the first cipher in its list that matches any one of the ciphers in the client's list. When the user verifies the preconfigured policies supported by ELB, the policy ??ELBSecurity Policy-2014-01?? supports server order preference.

#### NEW QUESTION 163

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 166

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 171

A user has launched an EC2 Windows instance from an instance store backed AMI. The user has also set the Instance initiated shutdown behavior to stop. What will happen when the user shuts down the OS?

- A. It will not allow the user to shutdown the OS when the shutdown behaviour is set to Stop
- B. It is not possible to set the termination behaviour to Stop for an Instance store backed AMI instance
- C. The instance will stay running but the OS will be shutdown
- D. The instance will be terminated

**Answer:** B

#### Explanation:

When the EC2 instance is launched from an instance store backed AMI, it will not allow the user to configure the shutdown behaviour to ??Stop??. It gives a warning that the instance does not have the EBS root volume.

#### NEW QUESTION 172

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

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

**Answer:** B

#### Explanation:

AWS S3 supports client side or server side encryption to encrypt all data at rest. The server side encryption can either have the S3 supplied AES-256 encryption key or the user can send the key along with each API call to supply his own encryption key (SSE-C.. If the bucket is versioning-enabled, each object version

uploaded by the user using the SSE-C feature can have its own encryption key. The user is responsible for tracking which encryption key was used for which object's version

#### NEW QUESTION 177

A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24. What will happen in this scenario?

- A. The VPC will modify the first subnet CIDR automatically to allow the second subnet IP range
- B. It is not possible to create a subnet with the same CIDR as VPC
- C. The second subnet will be created
- D. It will throw a CIDR overlaps error

**Answer:** D

#### Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet.

#### NEW QUESTION 182

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 185

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 190

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

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

**Answer:** D

#### Explanation:

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

#### NEW QUESTION 191

A user is receiving a notification from the RDS DB whenever there is a change in the DB security group. The user does not want to receive these notifications for only a month. Thus, he does not want to delete the notification. How can the user configure this?

- A. Change the Disable button for notification to `Yes` in the RDS console
- B. Set the send mail flag to false in the DB event notification console
- C. The only option is to delete the notification from the console
- D. Change the Enable button for notification to `No` in the RDS console

**Answer:** D

**Explanation:**

Amazon RDS uses the Amazon Simple Notification Service to provide a notification when an Amazon RDS event occurs. Event notifications are sent to the addresses that the user has provided while creating the subscription. The user can easily turn off the notification without deleting a subscription by setting the Enabled radio button to No in the Amazon RDS console or by setting the Enabled parameter to false using the CLI or Amazon RDS API.

**NEW QUESTION 194**

A user is trying to understand the CloudWatch metrics for the AWS services. It is required that the user should first understand the namespace for the AWS services. Which of the below mentioned is not a valid namespace for the AWS services?

- A. AWS/StorageGateway
- B. AWS/CloudTrail
- C. AWS/ElastiCache
- D. AWS/SWF

**Answer: B**

**Explanation:**

Amazon CloudWatch is basically a metrics repository. The AWS product puts metrics into this repository, and the user can retrieve the data or statistics based on those metrics. To distinguish the data for each service, the CloudWatch metric has a namespace. Namespaces are containers for metrics. All AWS services that provide the Amazon CloudWatch data use a namespace string, beginning with "AWS/". All the services which are supported by CloudWatch will have some namespace. CloudWatch does not monitor CloudTrail. Thus, the namespace ??AWS/CloudTrail?? is incorrect.

**NEW QUESTION 195**

A user has launched an EC2 instance store backed instance in the US-East-1a zone. The user created AMI #1 and copied it to the Europe region. After that, the user made a few updates to the application running in the US-East-1a zone. The user makes an AMI#2 after the changes. If the user launches a new instance in Europe from the AMI #1 copy, which of the below mentioned statements is true?

- A. The new instance will have the changes made after the AMI copy as AWS just copies the reference of the original AMI during the copyin
- B. Thus, the copied AMI will have all the updated data
- C. The new instance will have the changes made after the AMI copy since AWS keeps updating the AMI
- D. It is not possible to copy the instance store backed AMI from one region to another
- E. The new instance in the EU region will not have the changes made after the AMI copy

**Answer: D**

**Explanation:**

Within EC2, when the user copies an AMI, the new AMI is fully independent of the source AMI; there is no link to the original (source) AMI. The user can modify the source AMI without affecting the new AMI and vice versa. Therefore, in this case even if the source AMI is modified, the copied AMI of the EU region will not have the changes. Thus, after copy the user needs to copy the new source AMI to the destination region to get those changes.

**NEW QUESTION 200**

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 202**

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 203**

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 207**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 209**

A user has provisioned 2000 IOPS to the EBS volume. The application hosted on that EBS is experiencing less IOPS than provisioned. Which of the below mentioned options does not affect the IOPS of the volume?

- A. The application does not have enough IO for the volume
- B. The instance is EBS optimized
- C. The EC2 instance has 10 Gigabit Network connectivity
- D. The volume size is too large

**Answer:** D

**Explanation:**

When the application does not experience the expected IOPS or throughput of the PIOPS EBS volume that was provisioned, the possible root cause could be that the EC2 bandwidth is the limiting factor and the instance might not be either EBS-optimized or might not have 10 Gigabit network connectivity. Another possible cause for not experiencing the expected IOPS could also be that the user is not driving enough I/O to the EBS volumes. The size of the volume may not affect IOPS.

**NEW QUESTION 213**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 216**

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

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

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 221

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 222

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 227

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 232

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 235

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 236**

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 237**

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

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

**Answer:** A

**Explanation:**

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The default maximums for each of the IAM entities is given below:

Groups per AWS account: 100 Users per AWS account: 5000 Roles per AWS account: 250

Number of groups per user: 10 (that is, one user can be part of these many groups).

**NEW QUESTION 242**

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 245**

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

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

**Answer:** D

**Explanation:**

AWS S3 supports client side or server side encryption to encrypt all data at Rest. The server side

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

#### NEW QUESTION 248

A user has created a VPC with public and private subnets using the VPC Wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24. Which of the below mentioned entries are required in the main route table to allow the instances in VPC to communicate with each other?

- A. Destination : 20.0.0.0/24 and Target : VPC
- B. Destination : 20.0.0.0/16 and Target : ALL
- C. Destination : 20.0.0.0/0 and Target : ALL
- D. Destination : 20.0.0.0/24 and Target : Local

**Answer: D**

#### NEW QUESTION 253

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

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

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

**Answer: B**

#### NEW QUESTION 257

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

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

**Answer: A**

#### Explanation:

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

#### NEW QUESTION 259

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 264

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

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

**Answer: A**

#### NEW QUESTION 265

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

- A. AWS Direct Connect
- B. Placement Groups

- C. VPC private subnets
- D. EC2 Dedicated Instances
- E. Multiple Availability Zones

**Answer:** B

**Explanation:**

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

**NEW QUESTION 270**

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 272**

You have private video content in S3 that you want to serve to subscribed users on the Internet. User IDs, credentials, and subscriptions are stored in an Amazon RDS database.

Which configuration will allow you to securely serve private content to your users?

- A. Generate pre-signed URLs for each user as they request access to protected S3 content
- B. Create an IAM user for each subscribed user and assign the GetObject permission to each IAM user
- C. Create an S3 bucket policy that limits access to your private content to only your subscribed users' credentials
- D. Create a CloudFront Origin Identity user for your subscribed users and assign the GetObject permission to this user

**Answer:** D

**Explanation:**

Reference:

<https://java.awsblog.com/post/Tx1VE22EWFR4H86/Accessing-Private-Content-in-Amazon-CloudFront>

**NEW QUESTION 273**

Your mission is to create a lights-out datacenter environment, and you plan to use AWS OpsWorks to accomplish this. First you created a stack and added an App Server layer with an instance running in it. Next you added an application to the instance, and now you need to deploy a MySQL RDS database instance.

Which of the following answers accurately describe how to add a backend database server to an OpsWorks stack? Choose 3 answers

- A. Add a new database layer and then add recipes to the deploy actions of the database and App Server layers.
- B. Use OpsWorks' "Clone Stack" feature to create a second RDS stack in another Availability Zone for redundancy in the event of a failure in the Primary A
- C. To switch to the secondary RDS instance, set the [:database] attributes to values that are appropriate for your server which you can do by using custom JSON.
- D. The variables that characterize the RDS database connection?Xhost, user, and so on?Xare set using the corresponding values from the deploy JSON's [:deploy][:app\_name][:database] attributes.
- E. Cookbook attributes are stored in a repository, so OpsWorks requires that the "password": "your\_password" attribute for the RDS instance must be encrypted using at least a 256-bit key.
- F. Set up the connection between the app server and the RDS layer by using a custom recip
- G. The recipe configures the app server as required, typically by creating a configuration fil
- H. The recipe gets the connection data such as the host and database name from a set of attributes in the stack configuration and deployment JSON that AWS OpsWorks installs on every instance.

**Answer:** BCE

**NEW QUESTION 278**

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

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

**Answer:** B

**Explanation:**

[http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_ACLs.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html)

**NEW QUESTION 281**

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 285**

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 287**

An Auto-Scaling group spans 3 AZs and currently has 4 running EC2 instances. When Auto Scaling needs to terminate an EC2 instance by default, AutoScaling will:

Choose 2 answers

- A. Allow at least five minutes for Windows/Linux shutdown scripts to complete, before terminating the instance.
- B. Terminate the instance with the least active network connection
- C. If multiple instances meet this criterion, one will be randomly selected.
- D. Send an SNS notification, if configured to do so.
- E. Terminate an instance in the AZ which currently has 2 running EC2 instances.
- F. Randomly select one of the 3 AZs, and then terminate an instance in that AZ.

**Answer:** CD

**Explanation:**

<http://docs.aws.amazon.com/autoscaling/latest/userguide/as-instance-termination.html>

**NEW QUESTION 289**

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

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

**Answer:** C

**NEW QUESTION 290**

A syslog Administrator is created additional Amazon EC2 instances and receive an InstanceLimitExceeded error. What is the cause of the issue and how can it be resolve?

- A. The Administrator has requested too many instances at once and must request fewer instances in batches
- B. The concurrent running instance limit has been reached and an EC2 limit increase request must be filed with AWS Support
- C. AWS does not currently have enough available capacity and a different instance type must be used
- D. The Administrator must specify the maximum number of instances to be created provisioning EC instances

**Answer:** B

**Explanation:**

EC2 Service Limits: AWS sets limits for these resources on a per-region basis.

If you are getting an InstanceLimitExceeded error when you try to launch an instance, you have reached your concurrent running instance limit. For new AWS accounts, the default limit is 20. If you need additional running instances, complete the form at Request to Increase Amazon EC2 Instance Limit.

By default, all AWS accounts have a limit of 20 running instances at any time per region. If you attempt to start another one, even if it already existed in the stopped state, you will receive this error message.

To resolve this issue, you can do any of the following: Stop one of your other running instances

Contact AWS support and request your running EC2 instances quota limit be raised.

**NEW QUESTION 292**

A SysOps Administrator is asked to create an Amazon VPC IPv4 subnet that will support a minimum of 30 network resources simultaneously. What is the minimum CIDR netmask that will sustain this requirement?

- A. /25
- B. /26
- C. /27
- D. /28

**Answer:** C

**Explanation:**

### CIDR Available Hosts

The formula to calculate the number of assignable IP address to CIDR networks is similar to classful networking. Subtract the number of network bits from 32. Raise 2 to that power and subtract 2 for the network and broadcast addresses. For example, a /24 network has  $2^{32-24} - 2$  addresses available for host assignment.

CIDR Notation	Host Formula	Available Hosts
/8	$2^{32-8} - 2$	16,777,214
/9	$2^{32-9} - 2$	8,388,606
/10	$2^{32-10} - 2$	4,194,302
/11	$2^{32-11} - 2$	2,097,150
/12	$2^{32-12} - 2$	1,048,574
/13	$2^{32-13} - 2$	524,286
/14	$2^{32-14} - 2$	262,142
/15	$2^{32-15} - 2$	131,070
/16	$2^{32-16} - 2$	65,534
/17	$2^{32-17} - 2$	32,766
/18	$2^{32-18} - 2$	16,382
/19	$2^{32-19} - 2$	8,190
/20	$2^{32-20} - 2$	4,094
/21	$2^{32-21} - 2$	2,046
/22	$2^{32-22} - 2$	1,022
/23	$2^{32-23} - 2$	510
/24	$2^{32-24} - 2$	254
/25	$2^{32-25} - 2$	126
/26	$2^{32-26} - 2$	62
/27	$2^{32-27} - 2$	30
/28	$2^{32-28} - 2$	14
/29	$2^{32-29} - 2$	6
/30	$2^{32-30} - 2$	2

### NEW QUESTION 296

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 299

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 302

A company three-tier web application is not performing as well as expected. A manager has asked a System Administrator to analyser all the system involved and identity where the performance bottleneck exist. Which AWS service can be help find bottleneck?


- A. Analyse AWS CloudTrail logs to see which API call are taking the longest to execute
- B. Run a performance trace using Amazon Inspector to measure response tone between various API calls
- C. Create a rule in AWS Config to send an alert when the performance s noncompliant for each of the tiers
- D. Create an Amazon CloudWatch dashboard that contains Amazon EC2 and Amazon RDS metrics

**Answer: D**

**Explanation:**

Check the CloudWatch Latency metric

The Latency metric represents the time elapsed, in seconds, after the request leaves the load balancer until a response is received by the load balancer from a registered instance. The preferred statistic for this metric is average, which reports average latency for all requests. A high Latency average value typically indicates a problem with the backend server(s) rather than a problem with the load balancer. Check the maximum statistic to determine the number of latency data points that reach or exceed the load balancer idle timeout value. When latency data points meet or exceed the idle timeout value, it is likely that some requests are timing out, which initiates an HTTP 504 response to clients.



**NEW QUESTION 307**

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 309**

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

```
2 223342798652 eni-abc77deb 10.1.1.200 10.100.1.10 1123 25 17 70 48252 1515534437 1515535037 REJECT OK
```

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

**Answer: D**

**NEW QUESTION 314**

An errant process is known to use in an entire processor and run at 100%. A SysOps Administrator wants to automate restarting the instance once the problem occurs for more than minutes.

How can this be accomplished?

- A. Create an Amazon CloudWatch alarm on the Amazon EC2 instance with basic monitoring Enable an action to restart the instance
- B. Create a CloudWatch alarm for the EC2 instance with detailed monitoring Enable an action to restart the instance
- C. Create an AWS Lambda function to restart the EC2 instance triggered on a scheduled basis every 2 minutes
- D. Create a Lambda function start the EC2 instance triggered by EC2 health

**Answer: D**

**Explanation:**

You can use CloudWatch Events to trigger an AWS Lambda function to start and stop your EC2 instances at scheduled intervals.

Note: This article provides an example for a simple solution. For a more robust solution, see AWS Instance Scheduler.

Resolution

CloudWatch Events allows you to create an event that is triggered at a specified time or interval in response to events that take place in your account. For example, you can create an event using CloudWatch Events for a specific time of day, or you can create an alarm when CPU utilization for an instance reaches a specific threshold. You can also configure a Lambda function to start and stop instances when triggered by these events.

In this example, we use Lambda functions to start and stop EC2 instances, and then we use CloudWatch Events to start instances in the morning and stop the

instances at night.

1. Open the AWS Lambda console, and choose Create function.
2. Choose Author from scratch.
3. Enter a Name for your function, such as "StopEC2Instances."
4. From the Runtime drop-down menu, choose Python2.7.
5. Expand the Role drop-down menu, and then choose Create a custom role. This opens a new tab or window in your browser.
6. In the IAM Role drop-down menu, choose Create a new IAM Role, and enter a Role Name, such as "lambda\_start\_stop\_ec2."
7. Expand View Policy Document, choose Edit, and then choose Ok when prompted to read the documentation.

**NEW QUESTION 317**

A company is running an Oracle database engine that handles heavy online transaction processing (OLTP) structured data traffic. How can a SysOps administrator ensure that the database has high availability?

- A. Use Amazon DynamoDB to store the data
- B. Use Amazon RDS Multi-AZ deployment to store the data
- C. Use Amazon RDS read replicas in a different region to store the data
- D. Use an Amazon Redshift cluster to store the data

**Answer: B**

**NEW QUESTION 321**

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 324**

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 328**

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	Protocol	Port Range	Source	Destination	Comments
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.
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 329

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

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

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

**Answer:** ABD

**Explanation:**

[https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/viewing\\_metrics\\_with\\_cloudwatch.html](https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/viewing_metrics_with_cloudwatch.html)

#### NEW QUESTION 333

A SysOps Administrator is developing a cost-effective solution assist the Finance department with batch processing Fiancé is flexible on the batch processing schedule. and as long as the batch process runs within the same week, schedule interrupted are acceptable.

Which compute strategy can the SysOps Administrator use to meet the requirement?

- A. Amazon EC2 Reserved Instances
- B. Amazon EC2 Spot Instances
- C. Amazon EC2 Defeated Hosts
- D. Amazon EC2 On-Demand Instances

**Answer:** D

#### NEW QUESTION 335

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

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

**Answer:** D

**Explanation:**

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

#### NEW QUESTION 338

.....

## **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](#)**