



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

NEW QUESTION 2

You have identified network throughput as a bottleneck on your m1.small EC2 instance when uploading data into Amazon S3 in the same region.

How do you remedy this situation?

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 3

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

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

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

Answer: B

Explanation:

As the EC2 limit per region is max 20. You will need to fill an Amazon EC2 instance request form to increase the EC2 instances to 175.

http://aws.amazon.com/ec2/faqs/#How_many_instances_can_I_run_in_Amazon_EC2

I don't think the answer can be D, as the question says "expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks". To pre-warm your ELB, you have to put in a request to AWS. You can't do it.

Q: How do I reserve capacity for an existing, running instance?

To reserve capacity for a running instance, you can purchase a Reserved Instance or modify an existing reservation so it matches your instance's specifications.

You can purchase Reserved Instances via the Amazon EC2 Console or by using the `PurchaseReservedInstancesOffering` API. You can modify existing Reserved Instances via the Amazon EC2 Console or by using the `ModifyReservedInstances` API call.

In both cases, the reservation must match the following attributes of the running instance you want to cover:

Availability Zone (e.g., us-east-1a) Instance type (e.g., m3.large)

Platform (e.g., Linux/UNIX (Amazon VPC)) Tenancy (e.g., default)

Q: How do I control which instances are billed at the lower rate?

The `RunInstances` API command does not distinguish between On-Demand instances and the reservations that can be applied to them. When computing your bill, our system will automatically optimize which instances are charged at the lower rate to ensure you always pay the lowest amount. For information about hourly billing, and how it applies to Reserved Instances, see [Billing Benefits and Payment Options](#).

Q: How many Reserved Instances can I purchase?

You can purchase up to 20 Reserved Instances per Availability Zone each month. If you need additional Reserved Instances, complete the form found [here](#).

Information about previous generation Reserved Instance types can be found [here](#).

Q: Can I reassign my Reserved Instance from one instance type (e.g., c1.xlarge) to another (e.g., m1.large)?

No. A Reserved Instance is associated with a specific instance type for the duration of its term; however, you can change from one instance size (e.g., c3.large) to another (e.g., c3.xlarge) in the same type, if it is a Linux/UNIX Reserved Instance.

Q: Can I move a Reserved Instance from one region to another?

No. A Reserved Instance is associated with a specific region, which is fixed for the duration of the reservation's term.

Q: Can I modify a Reserved Instance?

Yes. You can request to modify active reservations that you own in one of the following ways: Move between Availability Zones within the same region.

Change the network platform from EC2-Classic to EC2-VPC (for EC2-Classic-enabled customers). Change the instance type of your Linux/UNIX Reserved Instances to a larger or smaller size in the same instance type (e.g., convert 8 m1.smalls into 4 m1.mediums, or vice versa).

Instance type modifications are only supported for Linux/UNIX platform reservations. However, due to licensing differences Linux Reserved Instances cannot be modified to RedHat or SUSE Linux Reserved Instances.

The reservations that you modify must have been purchased on the same day, be the same instance type, and in the same Availability Zone and region. It is not possible to combine reservations. However, if you have multiple instances in the same reservation (i.e., the reservation was purchased to apply to 10 instances), you can modify each of these instances either individually or as a whole.

Q: How do I request changes or modifications?

You can submit a modification request from the Amazon EC2 Console or by using the `ModifyReservedInstances` API. We process your requests as soon as

possible, depending on available capacity. There is no additional cost for modifying your Reserved Instances. To learn more about modification, see the Amazon EC2 User Guide.

NEW QUESTION 4

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 5

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

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

Answer: A

Explanation:

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

NEW QUESTION 6

You have set up individual AWS accounts for each project. You have been asked to make sure your AWS Infrastructure costs do not exceed the budget set per project for each month. Which of the following approaches can help ensure that you do not exceed the budget each month?

- A. Consolidate your accounts so you have a single bill for all accounts and projects
- B. Set up auto scaling with CloudWatch alarms using SNS to notify you when you are running too many instances in a given account
- C. Set up CloudWatch billing alerts for all AWS resources used by each project, with a notification occurring when the amount for each resource tagged to a particular project matches the budget allocated to the project.
- D. Set up CloudWatch billing alerts for all AWS resources used by each account, with email notifications when it hits 50%, 80% and 90% of its budgeted monthly spend

Answer: C

NEW QUESTION 7

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 8

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

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

Answer: A

Explanation:

See: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html#instance-store-lifetime>
However, data in the instance store is lost under the following circumstances:
?V The underlying disk drive fails

?V The instance stops
?V The instance terminates

NEW QUESTION 9

You have a server with a 500GB Amazon EBS data volume. The volume is 80% full. You need to back up the volume at regular intervals and be able to re-create the volume in a new Availability Zone in the shortest time possible. All applications using the volume can be paused for a period of a few minutes with no discernible user impact.

Which of the following backup methods will best fulfill your requirements?

- A. Take periodic snapshots of the EBS volume
- B. Use a third party Incremental backup application to back up to Amazon Glacier
- C. Periodically back up all data to a single compressed archive and archive to Amazon S3 using a parallelized multi-part upload
- D. Create another EBS volume in the second Availability Zone attach it to the Amazon EC2 instance, and use a disk manager to mirror the two disks

Answer: A

Explanation:

Since an EBS volume should be in the same AZ as the EC2 instance. You cannot connect a EBS volume in another AZ.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-restoring-volume.html> EBS volumes can only be attached to EC2 instances within the same Availability Zone.

NEW QUESTION 10

An organization is planning to create 5 different AWS accounts considering various security requirements. The organization wants to use a single payee account by using the consolidated billing option. Which of the below mentioned statements is true with respect to the above information?

- A. Master (Payee) account will get only the total bill and cannot see the cost incurred by each account
- B. Master (Payee) account can view only the AWS billing details of the linked accounts
- C. It is not recommended to use consolidated billing since the payee account will have access to the linked accounts
- D. Each AWS account needs to create an AWS billing policy to provide permission to the payee account

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. Consolidated billing enables the organization to see a combined view of the AWS charges incurred by each account as well as obtain a detailed cost report for each of the individual AWS accounts associated with the paying account. The payee account will not have any other access than billing data of linked accounts.

NEW QUESTION 10

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 13

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

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

Answer: A

Explanation:

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

NEW QUESTION 18

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 19

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 24

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

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

Answer: D

Explanation:

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

NEW QUESTION 27

You are building an online store on AWS that uses SQS to process your customer orders. Your backend system needs those messages in the same sequence the customer orders have been put in. How can you achieve that?

- A. It is not possible to do this with SQS
- B. You can use sequencing information on each message
- C. You can do this with SQS but you also need to use SWF
- D. Messages will arrive in the same order by default

Answer: B

Explanation:

Amazon SQS is engineered to always be available and deliver messages. One of the resulting tradeoffs is that SQS does not guarantee first in, first out delivery of messages. For many distributed applications, each message can stand on its own, and as long as all messages are delivered, the order is not important. If your system requires that order be preserved, you can place sequencing information in each message, so that you can reorder the messages when the queue returns them.

NEW QUESTION 32

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

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

Answer: D

Explanation:

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

NEW QUESTION 33

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 35

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 38

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

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

Answer: D

Explanation:

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

NEW QUESTION 40

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 42

A user has launched an EBS backed EC2 instance. What will be the difference while performing the restart or stop/start options on that instance?

- A. For restart it does not charge for an extra hour, while every stop/start it will be charged as a separate hour
- B. Every restart is charged by AWS as a separate hour, while multiple start/stop actions during a single hour will be counted as a single hour
- C. For every restart or start/stop it will be charged as a separate hour
- D. For restart it charges extra only once, while for every stop/start it will be charged as a separate hour

Answer: A

Explanation:

For an EC2 instance launched with an EBS backed AMI, each time the instance state is changed from stop to start/ running, AWS charges a full instance hour, even if these transitions happen multiple times within a single hour. Anyway, rebooting an instance AWS does not charge a new instance billing hour.

NEW QUESTION 47

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

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

Answer: C

Explanation:

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

NEW QUESTION 51

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

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

Answer: A

Explanation:

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

NEW QUESTION 55

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

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

Answer: D

Explanation:

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

NEW QUESTION 60

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

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

Answer: C

Explanation:

Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. The user can specify any date in the future to scale up or down during that period. As per Auto Scaling the user can schedule an action for up to a month in the future. Thus, it is recommended to wait until end of November before scheduling for Christmas.

NEW QUESTION 61

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 63

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

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

Answer: D

Explanation:

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

NEW QUESTION 68

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

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

Answer: D

Explanation:

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

NEW QUESTION 69

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

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

Answer: A

Explanation:

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

NEW QUESTION 70

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

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

Answer: B

Explanation:

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

NEW QUESTION 71

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 76

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

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

Answer: D

Explanation:

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

NEW QUESTION 79

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 83

A sysadmin has created the below mentioned policy on an S3 bucket named cloudacademy. What does this policy define?

```
"Statement": [{
  "Sid": "Stmt1388811069831",
  "Effect": "Allow", "Principal": { "AWS": "*" },
  "Action": [ "s3:GetObjectAcl", "s3:ListBucket"], "Resource": [ "arn:aws:s3:::cloudacademy"]
}]
```

- A. It will make the cloudacademy bucket as well as all its objects as public
- B. It will allow everyone to view the ACL of the bucket
- C. It will give an error as no object is defined as part of the policy while the action defines the rule about the object
- D. It will make the cloudacademy bucket as public

Answer: D

Explanation:

A sysadmin can grant permission to the S3 objects or the buckets to any user or make objects public using the bucket policy and user policy. Both use the JSON-based access policy language. Generally if 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. In the sample policy the action says `s3:ListBucket` for effect Allow on Resource `arn:aws:s3:::cloudacademy`. This will make the cloudacademy bucket public.

```
"Statement": [{
  "Sid": "Stmt1388811069831",
  "Effect": "Allow", "Principal": { "AWS": "*" },
  "Action": [ "s3:GetObjectAcl", "s3:ListBucket"], "Resource": [ "arn:aws:s3:::cloudacademy"]
}]
```

NEW QUESTION 84

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

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

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

Answer: D

Explanation:

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

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

```
"Action": "autoscaling:Describe*", "Resource": "*"
}
]
}
```

NEW QUESTION 87

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 89

A user has configured ELB with a TCP listener at ELB as well as on the back-end instances. The user wants to enable a proxy protocol to capture the source and destination IP information in the header. Which of the below mentioned statements helps the user understand a proxy protocol with TCP configuration?

- A. If the end user is requesting behind a proxy server then the user should not enable a proxy protocol on ELB
- B. ELB does not support a proxy protocol when it is listening on both the load balancer and the back- end instances
- C. Whether the end user is requesting from a proxy server or directly, it does not make a difference for the proxy protocol
- D. If the end user is requesting behind the proxy then the user should add the `??isproxy??` flag to the ELB Configuration

Answer: A

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. If the end user is requesting from a Proxy Protocol enabled proxy server, then the ELB admin should not enable the Proxy Protocol on the load balancer. If the Proxy Protocol is enabled on both the proxy server and the load balancer, the load balancer will add another header to the request which already has a header from the proxy server. This duplication may result in errors.

NEW QUESTION 91

A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling AlarmNotification (which notifies Auto Scaling for CloudWatch alarms) for a while. What will Auto Scaling do during this period?

- A. AWS will not receive the alarms from CloudWatch
- B. AWS will receive the alarms but will not execute the Auto Scaling policy
- C. Auto Scaling will execute the policy but it will not launch the instances until the process is resumed
- D. It is not possible to suspend the AlarmNotification process

Answer: B

Explanation:

Auto Scaling performs various processes, such as Launch, Terminate Alarm Notification etc. The user can also suspend individual process. The AlarmNotification process type accepts notifications from the Amazon CloudWatch alarms that are associated with the Auto Scaling group. If the user suspends this process type, Auto Scaling will not automatically execute the scaling policies that would be triggered by the alarms.

NEW QUESTION 92

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 97

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 101

A user has launched an RDS MySQL DB with the Multi AZ feature. The user has scheduled the scaling of instance storage during maintenance window. What is the correct order of events during maintenance window?

Perform maintenance on standby Promote standby to primary

Perform maintenance on original primary Promote original master back as primary

- A. 1, 2, 3, 4
- B. 1, 2, 3
- C. 2, 3, 1, 4

Answer: B

Explanation:

Running MySQL on the RDS DB instance as a Multi-AZ deployment can help the user reduce the

impact of a maintenance event, as the Amazon will conduct maintenance by following the steps in the below mentioned order:

Perform maintenance on standby Promote standby to primary

Perform maintenance on original primary, which becomes the new standby.

NEW QUESTION 103

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. The user's data centre has CIDR 172.28.0.0/12. The user has also setup a NAT instance (i-123456) to allow traffic to the internet from the VPN subnet. Which of the below mentioned options is not a valid entry for the main route table in this scenario?

- A. Destination: 20.0.1.0/24 and Target: i-12345
- B. Destination: 0.0.0.0/0 and Target: i-12345
- C. Destination: 172.28.0.0/12 and Target: vgw-12345
- D. Destination: 20.0.0.0/16 and Target: local

Answer: A

Explanation:

The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, 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. If the user has setup a NAT instance to route all the internet requests then all requests to the internet should be routed to it. All requests to the organization's DC will be routed to the VPN gateway.

Here are the valid entries for the main route table in this scenario:

Destination: 0.0.0.0/0 & Target: i-12345 (To route all internet traffic to the NAT Instance).

Destination: 172.28.0.0/12 & Target: vgw-12345 (To route all the organization's data centre traffic to the VPN gateway).

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

NEW QUESTION 105

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

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

Answer: D

Explanation:

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

NEW QUESTION 109

A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services does not provide detailed monitoring with CloudWatch?

- A. AWS EMR
- B. AWS RDS
- C. AWS ELB
- D. AWS Route53

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. Services, such as RDS, EC2, Auto Scaling, ELB, and Route 53 can provide the monitoring data every minute.

NEW QUESTION 112

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

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

Answer: B

Explanation:

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

NEW QUESTION 113

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

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

Answer: C

Explanation:

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

NEW QUESTION 117

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 120

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 124

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

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

Answer: D

Explanation:

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

NEW QUESTION 125

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

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

Answer: D

Explanation:

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

NEW QUESTION 127

A user has configured ELB with SSL using a security policy for secure negotiation between the client and load balancer. Which of the below mentioned security policies is supported by ELB?

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

Answer: C

Explanation:

Elastic Load Balancing uses a Secure Socket Layer (SSL) negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. ELB supports two policies: Predefined Security Policy, which comes with predefined cipher and SSL protocols; Custom Security Policy, which allows the user to configure a policy.

NEW QUESTION 131

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 132

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 136

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

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

Answer: D

Explanation:

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

NEW QUESTION 137

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

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

Answer: A

Explanation:

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

NEW QUESTION 138

A user has launched an EC2 Windows instance from an instance store backed AMI. The user wants to convert the AMI to an EBS backed AMI. How can the user convert it?

- A. Attach an EBS volume to the instance and unbundle all the AMI bundled data inside the EBS
- B. A Windows based instance store backed AMI cannot be converted to an EBS backed AMI
- C. It is not possible to convert an instance store backed AMI to an EBS backed AMI
- D. Attach an EBS volume and use the copy command to copy all the ephemeral content to the EBS Volume

Answer: B

Explanation:

Generally when a user has launched an EC2 instance from an instance store backed AMI, it can be converted to an EBS backed AMI provided the user has attached the EBS volume to the instance and unbundles the AMI data to it. However, if the instance is a Windows instance, AWS does not allow this. In this case, since the instance is a Windows instance, the user cannot convert it to an EBS backed AMI.

NEW QUESTION 140

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

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

Answer: A

Explanation:

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

NEW QUESTION 143

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

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

Answer: B

Explanation:

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

NEW QUESTION 147

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

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

Answer: AB

NEW QUESTION 150

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

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

Answer: E

Explanation:

Reference:

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

NEW QUESTION 154

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

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 155

How can software determine the public and private IP addresses of the Amazon EC2 instance that it is running on?

- A. Query the local instance metadata.
- B. Query the appropriate Amazon CloudWatch metric.
- C. Query the local instance userdata.
- D. Use ipconfig or ifconfig command.

Answer: B

NEW QUESTION 160

The compliance department within your multi-national organization requires that all data for your customers that reside in the European Union (EU) must not leave the EU and also data for customers that reside in the US must not leave the US without explicit authorization.

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

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

Answer: C

NEW QUESTION 163

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

- A. Controlling physical access to compute resources
- B. Patch management on the EC2 instance's operating system
- C. Encryption of EBS (Elastic Block Storage) volumes
- D. Life-cycle management of IAM credentials
- E. Decommissioning storage devices
- F. Security Group and ACL (Access Control List) settings

Answer: BCDF

Explanation:

Decommissioning is AWS responsibility not Customer.

NEW QUESTION 167

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 recipe
- G. The recipe configures the app server as required, typically by creating a configuration file
- H. The recipe gets the connection data such as the host and database name from a set of attributes in the stack configuration and deployment JSON that AWS OpsWorks installs on every instance.

Answer: BCE

NEW QUESTION 171

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

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

Answer: B

Explanation:

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

NEW QUESTION 173

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 resolved?

- 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 EC2 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 174

A Developer reports that an Amazon EC2 instance has failed. The developer reports that all the data was stored on the root volume is now gone. What is the explanation for this issue?

- A. The instance was using an Amazon EBS root volume
- B. The instance was using Amazon S3 as the root volume
- C. The instance was using an instance store root volume
- D. The root volume with the data exists but needs to be re-attached

Answer: A

Explanation:

If your instance is ebs-backed, then you will not lose ebs root volume storage if you launched it with the "delete-on-termination" set to false. See the `ec2-run-instances` command for more information. Other ephemeral volumes will be lost when the instance is stopped/terminated. If it is not ebs-backed, you will lose the root data when you terminate the instance (you cannot "stop" instance store instances).

NEW QUESTION 179

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

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

Answer: D

NEW QUESTION 181

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 184

A SysOps Administrator has attempted to copy an Marketplace AMI an associated billing Product code that was shared another account. When the copy process is attempted, it fails.

What action can be taken to successfully copy the AMI to the target destination?

- A. Use an EC2 instance in the account by using the shared AMI and then created an AMI from the instance
- B. Launch an EC2 instance in the account by using the shared AMI and then create an AMI from the instance
- C. Use the AWS CLI with the `--nobillingProduct` flag to execute the copy and ignore the `billingProductcode`.
- D. Create a VPC peering connection between the source and target account to facilitate the AMI copy process.

Answer: D

NEW QUESTION 188

A SysOps Administrator supports a legacy application that is hardcoded to service `example.com`. The application has recently been moved to AWS. The external DNS are managed by a third-party provider. The Administrator has set up an internal domain for `example.com` and configured this record using Amazon Route. What solution offers the MOST efficient way to have instances in the same account resolve to the Route 53 service instead of the provider?

- A. Hardcode the name server record to the internal Route 53 IP address tor each instance
- B. Enable DNS resolution in the subnets as required
- C. Ensure that DNS resolution is enabled on the VPC
- D. Create an OS-specific hardcoded entry tor DNS resolution to the private URL

Answer: C

Explanation:

Using DNS with Your VPC

Domain Name System (DNS) is a standard by which names used on the Internet are resolved to their corresponding IP addresses. A DNS hostname is a name that uniquely and absolutely names a computer; it's composed of a host name and a domain name. DNS servers resolve DNS hostnames to their corresponding IP addresses.

Public IPv4 addresses enable communication over the Internet, while private IPv4 addresses enable communication within the network of the instance (either EC2-Classic or a VPC). For more information, see [IP Addressing in Your VPC](#).

We provide an Amazon DNS server. To use your own DNS server, create a new set of DHCP options for your VPC. For more information, see [DHCP Options Sets](#).

Contents

DNS Hostnames

DNS Support in Your VPC DNS Limits

Viewing DNS Hostnames for Your EC2 Instance Updating DNS Support for Your VPC

Using Private Hosted Zones

NEW QUESTION 193

A company has created a separate AWS account for all development work to protect the production environment in this development account, developers have permission to manipulate IAM policies and roles. Corporate policies require that developers and blocked from accessing some services. What is the BEST way to

grant the developers privileges in the development account while still complying with corporate policies?

- A. Create a service control policy in AWS Organizations and apply it to the development account
- B. Create a customer managed policy in IAM and apply it to all users within the development account
- C. Create a job function policy in IAM and apply it to all users within the development account
- D. Create an IAM policy and apply it in API Gateway to restrict the development account

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-service-control-policies-in-aws-organizations/>

NEW QUESTION 194

The Security team has decided that there will be no public internet access to HTTP (TCP port 80) because it is moving to HTTPS for all incoming web traffic. The team a SysOps Administrator to provide a report on any security groups that are not compliant. What should the SysOps Administrator do to provide near real-time compliance reporting?

- A. Enable AWS Trusted Advisor and show the security team that the Security groups unrestricted access check will alarm
- B. Schedule an AWS Lambda function to run hourly to scan and evaluate all security groups and send a report to the Security team
- C. Use AWS Config to enable the restricted-common ports rule and add port 80 to the parameters
- D. Use Amazon Inspector to evaluate the security groups during scans and send the completed reports to the Security team

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/trustedadvisor/best-practices/>

NEW QUESTION 195

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

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

The route table for the subnet in which this instance is launched

| Destination | Target |
|-------------|--------|
| 10.0.0.0/16 | local |

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

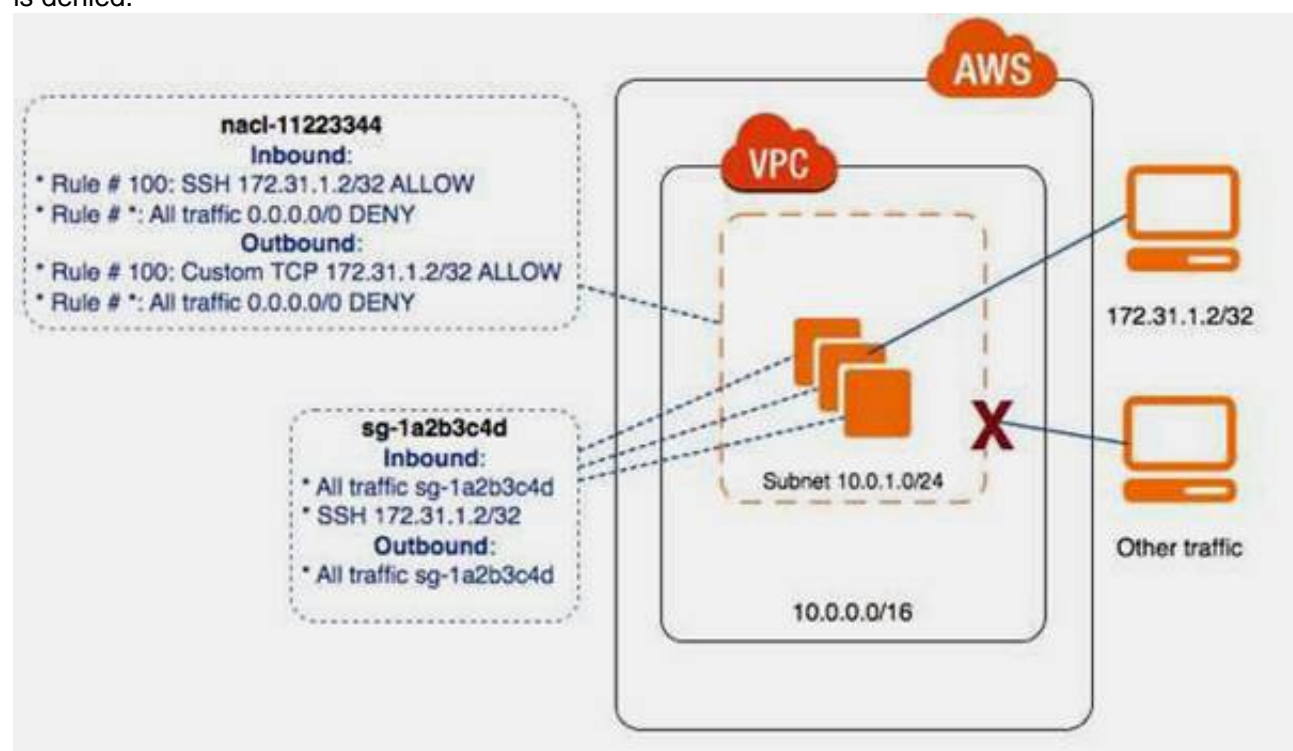
- A. The instance profile does not have explicit permissions to write objects to the S3 bucket.
- B. The route table does not have a rule for all traffic to pass through a NAT gateway.
- C. The route table does not have rule for all traffic to pass through an internet gateway

Answer: B

Explanation:

Controlling Access to Instances in a Subnet

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



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

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

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

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

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

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

This scenario gives you the flexibility to change the security groups or security group rules for your instances, and have the network ACL as the backup layer of defense. The network ACL rules apply to all instances in the subnet, so if you accidentally make your security group rules too permissive, the network ACL rules continue to permit access only from the single IP address. For example, the following rules are more permissive than the earlier rules ?X they allow inbound SSH access from any IP address.

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

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

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

NEW QUESTION 199

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 203

An Organization has been backing up their database backup to Amazon S3. A lifecycle rule has been created to transition these backups to Amazon Glacier storage class. The application development now to restore a backup.

Which step can an Administrator take to restore the backup to Amazon S3 storage?

- A. Create a new lifecycle rule to restore the backup from GLACIER storage class to Amazon S3 storage.
- B. Use the Amazon Glacier console to restore the backup from CLACIER storage class to Amazon S3 storage.
- C. Modify the existing lifecycle rule to restore the backup GKACIER storage class to Amazon S3 storage.
- D. Use the Amazon S3 console to restore the backup from CLACIER storage class to Amazon storage.

Answer: D

Explanation:

Restoring an Archived S3 Object

This topic explains how to use the Amazon S3 console to restore an object that has been archived to Glacier.

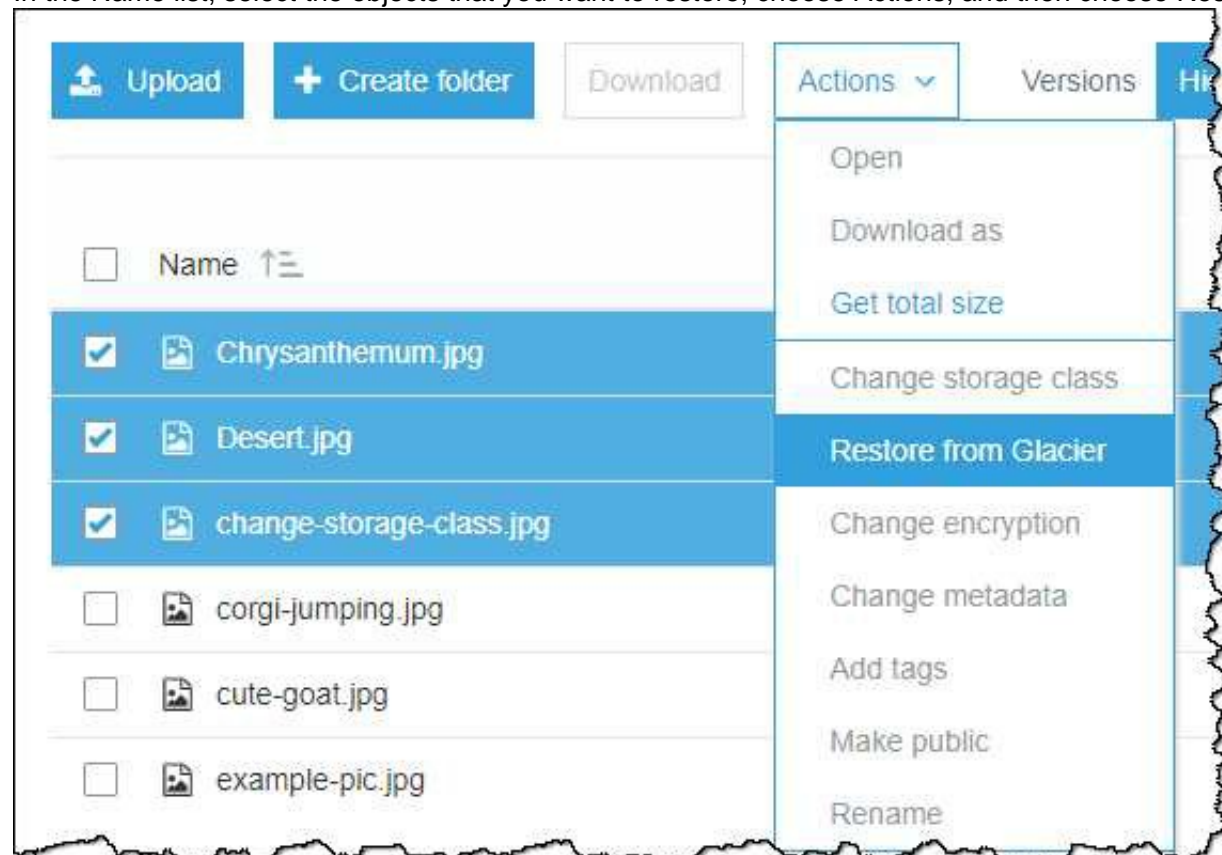
To restore archived S3 objects

Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.

In the Bucket name list, choose the name of the bucket that contains the objects that you want to restore.



In the Name list, select the objects that you want to restore, choose Actions, and then choose Restore from Glacier.



In the Initiate restore dialog box, type the number of days that you want your archived data to be accessible.

Choose one of the following retrieval options from the Retrieval options menu. Choose Bulk retrieval or Standard retrieval, and then choose Restore. Choose Expedited retrieval.

Restore objects from Glacier

Selection: 3 Objects, 0 Folders Total size: 1.7 MB Total objects: 3

Number of days the restored copy is available
 The restored copy in the Reduced Redundancy Storage (RRS) is automatically deleted after the specified number of days.

days
 Available until approximately 2018-12-05

Restore tier
 Glacier charges request fees and per GB retrieval fees, which vary based on the tier selected. See [S3 pricing](#)

☒ Bulk retrieval
 Typically within 5-12 hours

☐ Standard retrieval
 Typically within 3 - 5 hours

☐ Expedited retrieval
 Typically within 1 - 5 minutes when retrieving less than 250MB

If you have provisioned capacity, choose Restore to start a provisioned retrieval. If you have provisioned capacity, all of your expedited retrievals are served by your provisioned capacity. For more information about provisioned capacity, see [Provisioned Capacity](#).
 If you don't have provisioned capacity and you don't want to buy it, choose Restore.

If you don't have provisioned capacity, but you want to buy it, choose Add capacity unit, and then choose Buy. When you get the Purchase succeeded message, choose Restore to start provisioned retrieval.

☒ Expedited retrieval
Typically within 1 - 5 minutes when retrieving less than 250MB

Purchased capacity units: 0
Add 1 capacity unit

i Purchase 1 provisioned capacity unit. ✕

You will be immediately charged for each provisioned capacity unit and the purchase is not refundable. See [S3 pricing](#)

Provisioned capacity ensures that retrieval capacity for expedited retrievals is available when you need it. Each unit of capacity provides that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput.

Once purchased, provisioned capacity units will be available for your use in the current region for one month from the date of purchase.

Purchase

Cancel

Restore

NEW QUESTION 206

A company wants to send 70% of its inbound traffic to the us-east-1 region and 30% to the us-east region under normal; conditions. If all the servers go down in one of the regions, the company wants all the traffic to be re-routed to the other region.

- A. Configure an Application Load Balancer Target Group with weighted rules and a health check enabled
- B. Use a Network Load Balancer with sticky sessions enabled and weighted round robin with a 70/30 ratio
- C. Create two CNAME records in Amazon Route 53 enable dynamic traffic shaping with a 70/30 ratio
- D. Use a Route 53 weighted routing policy with a 70 /30 ratio and configure a health check

Answer: D

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-values-weighted-alias.html>

NEW QUESTION 209

A SysOps Administrator has an AWS Lambda function that performs maintenance on versions AWS resources. This function must be run nightly. Which is the MOST cost-effective solution?

- A. Launch a single I2.nano Amazon EC2 instance and create a Linux cron job to invoke the Lambda function at the same every right.
- B. Set up an Amazon CloudWatch metric alarm to invoke the Lambda function at the same time every night.
- C. Schedule a CloudWatch event to invoke the Lambda function at the same time every night.
- D. Implement a Chef recipe in Opsworks stack to invoke the Lambda function at the same time every night

Answer: C

Explanation:

Using AWS Lambda with Amazon CloudWatch Events

You can create a Lambda function and direct AWS Lambda to execute it on a regular schedule. You can specify a fixed rate (for example, execute a Lambda function every hour or 15 minutes), or you can specify a Cron expression. For more information on expressions schedules, see Schedule Expressions Using Rate or Cron.

This functionality is available when you create a Lambda function using the AWS Lambda console or the AWS CLI. To configure it using the AWS CLI, see Run an AWS Lambda Function on a Schedule Using the AWS CLI. The console provides CloudWatch Events as an event source. At the time of creating a Lambda function, you choose this event source and specify a time interval.

If you have made any manual changes to the permissions on your function, you may need to reapply the scheduled event access to your function. You can do that by using the following CLI command.

```
$ aws lambda add-permission --function-name function_name \
--action 'lambda:InvokeFunction' --principal events.amazonaws.com \
--statement-id 'statement_id' \
--source-arn arn:aws:events:region:account-id:rule/rule_name
```

Each AWS account can have up to 100 unique event sources of the CloudWatch Events- Schedule source type. Each of these can be the event source for up to five Lambda functions. That is, you can have up to 500 Lambda functions that can be executing on a schedule in your AWS account.

The console also provides a blueprint (lambda-canary) that uses the CloudWatch Events - Schedule source type. Using this blueprint, you can create a sample Lambda function and test this feature. The example code that the blueprint provides checks for the presence of a specific webpage and specific text string on the webpage. If either the webpage or the text string is not found, the Lambda function throws an error.

NEW QUESTION 214
.....

About ExamBible

[Your Partner of IT Exam](#)

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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

NEW QUESTION 2

You have identified network throughput as a bottleneck on your m1.small EC2 instance when uploading data into Amazon S3 in the same region.

How do you remedy this situation?

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 3

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

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

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

Answer: B

Explanation:

As the EC2 limit per region is max 20. You will need to fill an Amazon EC2 instance request form to increase the EC2 instances to 175.

http://aws.amazon.com/ec2/faqs/#How_many_instances_can_I_run_in_Amazon_EC2

I don't think the answer can be D, as the question says "expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks". To pre-warm your ELB, you have to put in a request to AWS. You can't do it.

Q: How do I reserve capacity for an existing, running instance?

To reserve capacity for a running instance, you can purchase a Reserved Instance or modify an existing reservation so it matches your instance's specifications.

You can purchase Reserved Instances via the Amazon EC2 Console or by using the `PurchaseReservedInstancesOffering` API. You can modify existing Reserved Instances via the Amazon EC2 Console or by using the `ModifyReservedInstances` API call.

In both cases, the reservation must match the following attributes of the running instance you want to cover:

Availability Zone (e.g., us-east-1a) Instance type (e.g., m3.large)

Platform (e.g., Linux/UNIX (Amazon VPC)) Tenancy (e.g., default)

Q: How do I control which instances are billed at the lower rate?

The `RunInstances` API command does not distinguish between On-Demand instances and the reservations that can be applied to them. When computing your bill, our system will automatically optimize which instances are charged at the lower rate to ensure you always pay the lowest amount. For information about hourly billing, and how it applies to Reserved Instances, see [Billing Benefits and Payment Options](#).

Q: How many Reserved Instances can I purchase?

You can purchase up to 20 Reserved Instances per Availability Zone each month. If you need additional Reserved Instances, complete the form found [here](#).

Information about previous generation Reserved Instance types can be found [here](#).

Q: Can I reassign my Reserved Instance from one instance type (e.g., c1.xlarge) to another (e.g., m1.large)?

No. A Reserved Instance is associated with a specific instance type for the duration of its term; however, you can change from one instance size (e.g., c3.large) to another (e.g., c3.xlarge) in the same type, if it is a Linux/UNIX Reserved Instance.

Q: Can I move a Reserved Instance from one region to another?

No. A Reserved Instance is associated with a specific region, which is fixed for the duration of the reservation's term.

Q: Can I modify a Reserved Instance?

Yes. You can request to modify active reservations that you own in one of the following ways: Move between Availability Zones within the same region.

Change the network platform from EC2-Classic to EC2-VPC (for EC2-Classic-enabled customers). Change the instance type of your Linux/UNIX Reserved Instances to a larger or smaller size in the same instance type (e.g., convert 8 m1.smalls into 4 m1.mediums, or vice versa).

Instance type modifications are only supported for Linux/UNIX platform reservations. However, due to licensing differences Linux Reserved Instances cannot be modified to RedHat or SUSE Linux Reserved Instances.

The reservations that you modify must have been purchased on the same day, be the same instance type, and in the same Availability Zone and region. It is not possible to combine reservations. However, if you have multiple instances in the same reservation (i.e., the reservation was purchased to apply to 10 instances), you can modify each of these instances either individually or as a whole.

Q: How do I request changes or modifications?

You can submit a modification request from the Amazon EC2 Console or by using the `ModifyReservedInstances` API. We process your requests as soon as

possible, depending on available capacity. There is no additional cost for modifying your Reserved Instances. To learn more about modification, see the Amazon EC2 User Guide.

NEW QUESTION 4

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 5

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

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

Answer: A

Explanation:

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

NEW QUESTION 6

You have set up individual AWS accounts for each project. You have been asked to make sure your AWS Infrastructure costs do not exceed the budget set per project for each month. Which of the following approaches can help ensure that you do not exceed the budget each month?

- A. Consolidate your accounts so you have a single bill for all accounts and projects
- B. Set up auto scaling with CloudWatch alarms using SNS to notify you when you are running too many instances in a given account
- C. Set up CloudWatch billing alerts for all AWS resources used by each project, with a notification occurring when the amount for each resource tagged to a particular project matches the budget allocated to the project.
- D. Set up CloudWatch billing alerts for all AWS resources used by each account, with email notifications when it hits 50%, 80% and 90% of its budgeted monthly spend

Answer: C

NEW QUESTION 7

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 8

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

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

Answer: A

Explanation:

See: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html#instance-store-lifetime>
However, data in the instance store is lost under the following circumstances:
?V The underlying disk drive fails

?V The instance stops
?V The instance terminates

NEW QUESTION 9

You have a server with a 500GB Amazon EBS data volume. The volume is 80% full. You need to back up the volume at regular intervals and be able to re-create the volume in a new Availability Zone in the shortest time possible. All applications using the volume can be paused for a period of a few minutes with no discernible user impact.

Which of the following backup methods will best fulfill your requirements?

- A. Take periodic snapshots of the EBS volume
- B. Use a third party Incremental backup application to back up to Amazon Glacier
- C. Periodically back up all data to a single compressed archive and archive to Amazon S3 using a parallelized multi-part upload
- D. Create another EBS volume in the second Availability Zone attach it to the Amazon EC2 instance, and use a disk manager to mirror the two disks

Answer: A

Explanation:

Since an EBS volume should be in the same AZ as the EC2 instance. You cannot connect a EBS volume in another AZ.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-restoring-volume.html> EBS volumes can only be attached to EC2 instances within the same Availability Zone.

NEW QUESTION 10

An organization is planning to create 5 different AWS accounts considering various security requirements. The organization wants to use a single payee account by using the consolidated billing option. Which of the below mentioned statements is true with respect to the above information?

- A. Master (Payee) account will get only the total bill and cannot see the cost incurred by each account
- B. Master (Payee) account can view only the AWS billing details of the linked accounts
- C. It is not recommended to use consolidated billing since the payee account will have access to the linked accounts
- D. Each AWS account needs to create an AWS billing policy to provide permission to the payee account

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. Consolidated billing enables the organization to see a combined view of the AWS charges incurred by each account as well as obtain a detailed cost report for each of the individual AWS accounts associated with the paying account. The payee account will not have any other access than billing data of linked accounts.

NEW QUESTION 10

A user has setup a CloudWatch alarm on an EC2 instance 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 13

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

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

Answer: A

Explanation:

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

NEW QUESTION 18

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 19

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 24

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

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

Answer: D

Explanation:

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

NEW QUESTION 27

You are building an online store on AWS that uses SQS to process your customer orders. Your backend system needs those messages in the same sequence the customer orders have been put in. How can you achieve that?

- A. It is not possible to do this with SQS
- B. You can use sequencing information on each message
- C. You can do this with SQS but you also need to use SWF
- D. Messages will arrive in the same order by default

Answer: B

Explanation:

Amazon SQS is engineered to always be available and deliver messages. One of the resulting tradeoffs is that SQS does not guarantee first in, first out delivery of messages. For many distributed applications, each message can stand on its own, and as long as all messages are delivered, the order is not important. If your system requires that order be preserved, you can place sequencing information in each message, so that you can reorder the messages when the queue returns them.

NEW QUESTION 32

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

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

Answer: D

Explanation:

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

NEW QUESTION 33

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 35

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 38

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

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

Answer: D

Explanation:

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

NEW QUESTION 40

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 42

A user has launched an EBS backed EC2 instance. What will be the difference while performing the restart or stop/start options on that instance?

- A. For restart it does not charge for an extra hour, while every stop/start it will be charged as a separate hour
- B. Every restart is charged by AWS as a separate hour, while multiple start/stop actions during a single hour will be counted as a single hour
- C. For every restart or start/stop it will be charged as a separate hour
- D. For restart it charges extra only once, while for every stop/start it will be charged as a separate hour

Answer: A

Explanation:

For an EC2 instance launched with an EBS backed AMI, each time the instance state is changed from stop to start/ running, AWS charges a full instance hour, even if these transitions happen multiple times within a single hour. Anyway, rebooting an instance AWS does not charge a new instance billing hour.

NEW QUESTION 47

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

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

Answer: C

Explanation:

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

NEW QUESTION 51

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

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

Answer: A

Explanation:

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

NEW QUESTION 55

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

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

Answer: D

Explanation:

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

NEW QUESTION 60

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

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

Answer: C

Explanation:

Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. The user can specify any date in the future to scale up or down during that period. As per Auto Scaling the user can schedule an action for up to a month in the future. Thus, it is recommended to wait until end of November before scheduling for Christmas.

NEW QUESTION 61

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 63

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

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

Answer: D

Explanation:

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

NEW QUESTION 68

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

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

Answer: D

Explanation:

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

NEW QUESTION 69

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

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

Answer: A

Explanation:

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

NEW QUESTION 70

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

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

Answer: B

Explanation:

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

NEW QUESTION 71

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 76

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

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

Answer: D

Explanation:

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

NEW QUESTION 79

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 83

A sysadmin has created the below mentioned policy on an S3 bucket named cloudacademy. What does this policy define?

```
"Statement": [{
  "Sid": "Stmt1388811069831",
  "Effect": "Allow", "Principal": { "AWS": "*" },
  "Action": [ "s3:GetObjectAcl", "s3:ListBucket"], "Resource": [ "arn:aws:s3:::cloudacademy"]
}]
```

- A. It will make the cloudacademy bucket as well as all its objects as public
- B. It will allow everyone to view the ACL of the bucket
- C. It will give an error as no object is defined as part of the policy while the action defines the rule about the object
- D. It will make the cloudacademy bucket as public

Answer: D

Explanation:

A sysadmin can grant permission to the S3 objects or the buckets to any user or make objects public using the bucket policy and user policy. Both use the JSON-based access policy language. Generally if 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. In the sample policy the action says `s3:ListBucket` for effect Allow on Resource `arn:aws:s3:::cloudacademy`. This will make the cloudacademy bucket public.

```
"Statement": [{
  "Sid": "Stmt1388811069831",
  "Effect": "Allow", "Principal": { "AWS": "*" },
  "Action": [ "s3:GetObjectAcl", "s3:ListBucket"], "Resource": [ "arn:aws:s3:::cloudacademy"]
}]
```

NEW QUESTION 84

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

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

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

Answer: D

Explanation:

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

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

```
"Action": "autoscaling:Describe*", "Resource": "*"
}
]
}
```

NEW QUESTION 87

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 89

A user has configured ELB with a TCP listener at ELB as well as on the back-end instances. The user wants to enable a proxy protocol to capture the source and destination IP information in the header. Which of the below mentioned statements helps the user understand a proxy protocol with TCP configuration?

- A. If the end user is requesting behind a proxy server then the user should not enable a proxy protocol on ELB
- B. ELB does not support a proxy protocol when it is listening on both the load balancer and the back- end instances
- C. Whether the end user is requesting from a proxy server or directly, it does not make a difference for the proxy protocol
- D. If the end user is requesting behind the proxy then the user should add the `??isproxy??` flag to the ELB Configuration

Answer: A

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. If the end user is requesting from a Proxy Protocol enabled proxy server, then the ELB admin should not enable the Proxy Protocol on the load balancer. If the Proxy Protocol is enabled on both the proxy server and the load balancer, the load balancer will add another header to the request which already has a header from the proxy server. This duplication may result in errors.

NEW QUESTION 91

A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling AlarmNotification (which notifies Auto Scaling for CloudWatch alarms) for a while. What will Auto Scaling do during this period?

- A. AWS will not receive the alarms from CloudWatch
- B. AWS will receive the alarms but will not execute the Auto Scaling policy
- C. Auto Scaling will execute the policy but it will not launch the instances until the process is resumed
- D. It is not possible to suspend the AlarmNotification process

Answer: B

Explanation:

Auto Scaling performs various processes, such as Launch, Terminate Alarm Notification etc. The user can also suspend individual process. The AlarmNotification process type accepts notifications from the Amazon CloudWatch alarms that are associated with the Auto Scaling group. If the user suspends this process type, Auto Scaling will not automatically execute the scaling policies that would be triggered by the alarms.

NEW QUESTION 92

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 97

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 101

A user has launched an RDS MySQL DB with the Multi AZ feature. The user has scheduled the scaling of instance storage during maintenance window. What is the correct order of events during maintenance window?

Perform maintenance on standby Promote standby to primary

Perform maintenance on original primary Promote original master back as primary

- A. 1, 2, 3, 4
- B. 1, 2, 3
- C. 2, 3, 1, 4

Answer: B

Explanation:

Running MySQL on the RDS DB instance as a Multi-AZ deployment can help the user reduce the impact of a maintenance event, as the Amazon will conduct maintenance by following the steps in the below mentioned order:

Perform maintenance on standby Promote standby to primary

Perform maintenance on original primary, which becomes the new standby.

NEW QUESTION 103

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. The user's data centre has CIDR 172.28.0.0/12. The user has also setup a NAT instance (i-123456) to allow traffic to the internet from the VPN subnet. Which of the below mentioned options is not a valid entry for the main route table in this scenario?

- A. Destination: 20.0.1.0/24 and Target: i-12345
- B. Destination: 0.0.0.0/0 and Target: i-12345
- C. Destination: 172.28.0.0/12 and Target: vgw-12345
- D. Destination: 20.0.0.0/16 and Target: local

Answer: A

Explanation:

The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, 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. If the user has setup a NAT instance to route all the internet requests then all requests to the internet should be routed to it. All requests to the organization's DC will be routed to the VPN gateway.

Here are the valid entries for the main route table in this scenario:

Destination: 0.0.0.0/0 & Target: i-12345 (To route all internet traffic to the NAT Instance).

Destination: 172.28.0.0/12 & Target: vgw-12345 (To route all the organization's data centre traffic to the VPN gateway).

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

NEW QUESTION 105

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

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

Answer: D

Explanation:

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

NEW QUESTION 109

A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services does not provide detailed monitoring with CloudWatch?

- A. AWS EMR
- B. AWS RDS
- C. AWS ELB
- D. AWS Route53

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. Services, such as RDS, EC2, Auto Scaling, ELB, and Route 53 can provide the monitoring data every minute.

NEW QUESTION 112

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

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

Answer: B

Explanation:

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

NEW QUESTION 113

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

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

Answer: C

Explanation:

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

NEW QUESTION 117

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 120

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 124

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

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

Answer: D

Explanation:

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

NEW QUESTION 125

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

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

Answer: D

Explanation:

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

NEW QUESTION 127

A user has configured ELB with SSL using a security policy for secure negotiation between the client and load balancer. Which of the below mentioned security policies is supported by ELB?

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

Answer: C

Explanation:

Elastic Load Balancing uses a Secure Socket Layer (SSL) negotiation configuration which is known as a Security Policy. It is used to negotiate the SSL connections between a client and the load balancer. ELB supports two policies: Predefined Security Policy, which comes with predefined cipher and SSL protocols; Custom Security Policy, which allows the user to configure a policy.

NEW QUESTION 131

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 132

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 136

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

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

Answer: D

Explanation:

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

NEW QUESTION 137

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

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

Answer: A

Explanation:

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

NEW QUESTION 138

A user has launched an EC2 Windows instance from an instance store backed AMI. The user wants to convert the AMI to an EBS backed AMI. How can the user convert it?

- A. Attach an EBS volume to the instance and unbundle all the AMI bundled data inside the EBS
- B. A Windows based instance store backed AMI cannot be converted to an EBS backed AMI
- C. It is not possible to convert an instance store backed AMI to an EBS backed AMI
- D. Attach an EBS volume and use the copy command to copy all the ephemeral content to the EBS Volume

Answer: B

Explanation:

Generally when a user has launched an EC2 instance from an instance store backed AMI, it can be converted to an EBS backed AMI provided the user has attached the EBS volume to the instance and unbundles the AMI data to it. However, if the instance is a Windows instance, AWS does not allow this. In this case, since the instance is a Windows instance, the user cannot convert it to an EBS backed AMI.

NEW QUESTION 140

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

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

Answer: A

Explanation:

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

NEW QUESTION 143

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

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

Answer: B

Explanation:

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

NEW QUESTION 147

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

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

Answer: AB

NEW QUESTION 150

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

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

Answer: E

Explanation:

Reference:

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

NEW QUESTION 154

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

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 155

How can software determine the public and private IP addresses of the Amazon EC2 instance that it is running on?

- A. Query the local instance metadata.
- B. Query the appropriate Amazon CloudWatch metric.
- C. Query the local instance userdata.
- D. Use ipconfig or ifconfig command.

Answer: B

NEW QUESTION 160

The compliance department within your multi-national organization requires that all data for your customers that reside in the European Union (EU) must not leave the EU and also data for customers that reside in the US must not leave the US without explicit authorization.

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

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

Answer: C

NEW QUESTION 163

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

- A. Controlling physical access to compute resources
- B. Patch management on the EC2 instance's operating system
- C. Encryption of EBS (Elastic Block Storage) volumes
- D. Life-cycle management of IAM credentials
- E. Decommissioning storage devices
- F. Security Group and ACL (Access Control List) settings

Answer: BCDF

Explanation:

Decommissioning is AWS responsibility not Customer.

NEW QUESTION 167

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 recipe
- G. The recipe configures the app server as required, typically by creating a configuration file
- H. The recipe gets the connection data such as the host and database name from a set of attributes in the stack configuration and deployment JSON that AWS OpsWorks installs on every instance.

Answer: BCE

NEW QUESTION 171

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

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

Answer: B

Explanation:

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

NEW QUESTION 173

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 resolved?

- 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 EC2 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 174

A Developer reports that an Amazon EC2 instance has failed. The developer reports that all the data that was stored on the root volume is now gone. What is the explanation for this issue?

- A. The instance was using an Amazon EBS root volume
- B. The instance was using Amazon S3 as the root volume
- C. The instance was using an instance store root volume
- D. The root volume with the data exists but needs to be re-attached

Answer: A

Explanation:

If your instance is ebs-backed, then you will not lose ebs root volume storage if you launched it with the "delete-on-termination" set to false. See the `ec2-run-instances` command for more information. Other ephemeral volumes will be lost when the instance is stopped/terminated. If it is not ebs-backed, you will lose the root data when you terminate the instance (you cannot "stop" instance store instances).

NEW QUESTION 179

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

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

Answer: D

NEW QUESTION 181

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 184

A SysOps Administrator has attempted to copy an Marketplace AMI an associated billing Product code that was shared another account. When the copy process is attempted, it fails.

What action can be taken to successfully copy the AMI to the target destination?

- A. Use an EC2 instance in the account by using the shared AMI and then created an AMI from the instance
- B. Launch an EC2 instance in the account by using the shared AMI and then create an AMI from the instance
- C. Use the AWS CLI with the `--nobillingProduct` flag to execute the copy and ignore the `billingProductcode`.
- D. Create a VPC peering connection between the source and target account to facilitate the AMI copy process.

Answer: D

NEW QUESTION 188

A SysOps Administrator supports a legacy application that is hardcoded to service `example.com`. The application has recently been moved to AWS. The external DNS are managed by a third-party provider. The Administrator has set up an internal domain for `example.com` and configured this record using Amazon Route. What solution offers the MOST efficient way to have instances in the same account resolve to the Route 53 service instead of the provider?

- A. Hardcode the name server record to the internal Route 53 IP address tor each instance
- B. Enable DNS resolution in the subnets as required
- C. Ensure that DNS resolution is enabled on the VPC
- D. Create an OS-specific hardcoded entry tor DNS resolution to the private URL

Answer: C

Explanation:

Using DNS with Your VPC

Domain Name System (DNS) is a standard by which names used on the Internet are resolved to their corresponding IP addresses. A DNS hostname is a name that uniquely and absolutely names a computer; it's composed of a host name and a domain name. DNS servers resolve DNS hostnames to their corresponding IP addresses.

Public IPv4 addresses enable communication over the Internet, while private IPv4 addresses enable communication within the network of the instance (either EC2-Classic or a VPC). For more information, see [IP Addressing in Your VPC](#).

We provide an Amazon DNS server. To use your own DNS server, create a new set of DHCP options for your VPC. For more information, see [DHCP Options Sets](#).

Contents

DNS Hostnames

DNS Support in Your VPC DNS Limits

Viewing DNS Hostnames for Your EC2 Instance Updating DNS Support for Your VPC

Using Private Hosted Zones

NEW QUESTION 193

A company has created a separate AWS account for all development work to protect the production environment in this development account, developers have permission to manipulate IAM policies and roles. Corporate policies require that developers and blocked from accessing some services. What is the BEST way to

grant the developers privileges in the development account while still complying with corporate policies?

- A. Create a service control policy in AWS Organizations and apply it to the development account
- B. Create a customer managed policy in IAM and apply it to all users within the development account
- C. Create a job function policy in IAM and apply it to all users within the development account
- D. Create an IAM policy and apply it in API Gateway to restrict the development account

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-service-control-policies-in-aws-organizations/>

NEW QUESTION 194

The Security team has decided that there will be no public internet access to HTTP (TCP port 80) because it is moving to HTTPS for all incoming web traffic. The team a SysOps Administrator to provide a report on any security groups that are not compliant. What should the SysOps Administrator do to provide near real-time compliance reporting?

- A. Enable AWS Trusted Advisor and show the security team that the Security groups unrestricted access check will alarm
- B. Schedule an AWS Lambda function to run hourly to scan and evaluate all security groups and send a report to the Security team
- C. Use AWS Config to enable the restricted-common ports rule and add port 80 to the parameters
- D. Use Amazon Inspector to evaluate the security groups during scans and send the completed reports to the Security team

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/trustedadvisor/best-practices/>

NEW QUESTION 195

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

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

The route table for the subnet in which this instance is launched

| Destination | Target |
|-------------|--------|
| 10.0.0.0/16 | local |

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


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

Answer: B

Explanation:

Controlling Access to Instances in a Subnet

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



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

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

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

| Rule # | Type | Protocol | Port Range | Source | Allow/Deny | Comments |
|--------|-------------|----------|------------|---------------|------------|---|
| 100 | SSH | TCP | 22 | 172.31.1.2/32 | ALLOW | Allows inbound traffic from the remote computer. If the instance is a Windows computer, then this rule must use the RDP protocol for port 3389 instead. |
| * | All traffic | All | All | 0.0.0.0/0 | DENY | Denies all other inbound traffic that does not match the previous rule. |
| Rule # | Type | Protocol | Port Range | Destination | Allow/Deny | Comments |
| 100 | Custom TCP | TCP | 1024-65535 | 172.31.1.2/32 | ALLOW | Allows outbound responses to the remote computer. Network ACLs are stateless, therefore this rule is required to allow response traffic for inbound requests. |
| * | All traffic | All | All | 0.0.0.0/0 | DENY | Denies all other outbound traffic that does not match the previous rule. |

This scenario gives you the flexibility to change the security groups or security group rules for your instances, and have the network ACL as the backup layer of defense. The network ACL rules apply to all instances in the subnet, so if you accidentally make your security group rules too permissive, the network ACL rules continue to permit access only from the single IP address. For example, the following rules are more permissive than the earlier rules ?X they allow inbound SSH access from any IP address.

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

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

NEW QUESTION 199

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 203

An Organization has been backing up their database backup to Amazon S3. A lifecycle rule has been created to transition these backups to Amazon Glacier storage class. The application development now to restore a backup.

Which step can an Administrator take to restore the backup to Amazon S3 storage?

- A. Create a new lifecycle rule to restore the backup from GLACIER storage class to Amazon S3 storage.
- B. Use the Amazon Glacier console to restore the backup from CLACIER storage class to Amazon S3 storage.
- C. Modify the existing lifecycle rule to restore the backup GKACIER storage class to Amazon S3 storage.
- D. Use the Amazon S3 console to restore the backup from CLACIER storage class to Amazon storage.

Answer: D

Explanation:

Restoring an Archived S3 Object

This topic explains how to use the Amazon S3 console to restore an object that has been archived to Glacier.

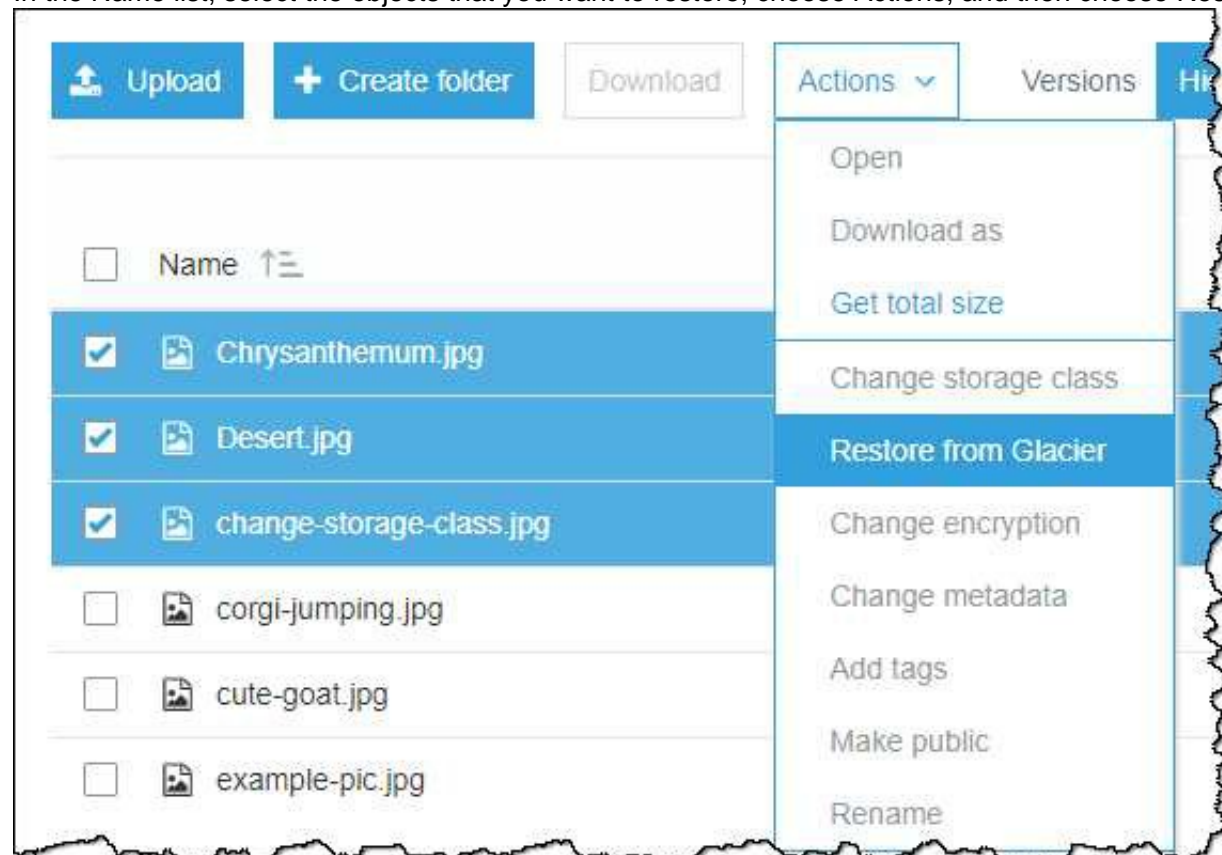
To restore archived S3 objects

Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.

In the Bucket name list, choose the name of the bucket that contains the objects that you want to restore.



In the Name list, select the objects that you want to restore, choose Actions, and then choose Restore from Glacier.



In the Initiate restore dialog box, type the number of days that you want your archived data to be accessible.

Choose one of the following retrieval options from the Retrieval options menu. Choose Bulk retrieval or Standard retrieval, and then choose Restore. Choose Expedited retrieval.

Restore objects from Glacier

Selection: 3 Objects, 0 Folders Total size: 1.7 MB Total objects: 3

Number of days the restored copy is available
 The restored copy in the Reduced Redundancy Storage (RRS) is automatically deleted after the specified number of days.

days
 Available until approximately 2018-12-05

Restore tier
 Glacier charges request fees and per GB retrieval fees, which vary based on the tier selected. See [S3 pricing](#)

☒ Bulk retrieval
 Typically within 5-12 hours

☐ Standard retrieval
 Typically within 3 - 5 hours

☐ Expedited retrieval
 Typically within 1 - 5 minutes when retrieving less than 250MB


If you have provisioned capacity, choose Restore to start a provisioned retrieval. If you have provisioned capacity, all of your expedited retrievals are served by your provisioned capacity. For more information about provisioned capacity, see [Provisioned Capacity](#).
 If you don't have provisioned capacity and you don't want to buy it, choose Restore.

If you don't have provisioned capacity, but you want to buy it, choose Add capacity unit, and then choose Buy. When you get the Purchase succeeded message, choose Restore to start provisioned retrieval.

☒ Expedited retrieval
Typically within 1 - 5 minutes when retrieving less than 250MB

Purchased capacity units: 0

i Purchase 1 provisioned capacity unit. ✕

You will be immediately charged for each provisioned capacity unit and the purchase is not refundable. See [S3 pricing](#) 

Provisioned capacity ensures that retrieval capacity for expedited retrievals is available when you need it. Each unit of capacity provides that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput.

Once purchased, provisioned capacity units will be available for your use in the current region for one month from the date of purchase.

NEW QUESTION 206

A company wants to send 70% of its inbound traffic to the us-east-1 region and 30% to the us-east region under normal; conditions. If all the servers go down in one of the regions, the company wants all the traffic to be re-routed to the other region.

- A. Configure an Application Load Balancer Target Group with weighted rules and a health check enabled
- B. Use a Network Load Balancer with sticky sessions enabled and weighted round robin with a 70/30 ratio
- C. Create two CNAME records in Amazon Route 53 enable dynamic traffic shaping with a 70/30 ratio
- D. Use a Route 53 weighted routing policy with a 70 /30 ratio and configure a health check

Answer: D

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-values-weighted-alias.html>

NEW QUESTION 209

A SysOps Administrator has an AWS Lambda function that performs maintenance on various AWS resources. This function must be run nightly. Which is the MOST cost-effective solution?

- A. Launch a single t2.nano Amazon EC2 instance and create a Linux cron job to invoke the Lambda function at the same time every night.
- B. Set up an Amazon CloudWatch metric alarm to invoke the Lambda function at the same time every night.
- C. Schedule a CloudWatch event to invoke the Lambda function at the same time every night.
- D. Implement a Chef recipe in Opsworks stack to invoke the Lambda function at the same time every night

Answer: C

Explanation:

Using AWS Lambda with Amazon CloudWatch Events

You can create a Lambda function and direct AWS Lambda to execute it on a regular schedule. You can specify a fixed rate (for example, execute a Lambda function every hour or 15 minutes), or you can specify a Cron expression. For more information on expressions schedules, see [Schedule Expressions Using Rate or Cron](#).

This functionality is available when you create a Lambda function using the AWS Lambda console or the AWS CLI. To configure it using the AWS CLI, see [Run an AWS Lambda Function on a Schedule Using the AWS CLI](#). The console provides CloudWatch Events as an event source. At the time of creating a Lambda function, you choose this event source and specify a time interval.

If you have made any manual changes to the permissions on your function, you may need to reapply the scheduled event access to your function. You can do that by using the following CLI command.

```
$ aws lambda add-permission --function-name function_name \
--action 'lambda:InvokeFunction' --principal events.amazonaws.com \
--statement-id 'statement_id' \
--source-arn arn:aws:events:region:account-id:rule/rule_name
```

Each AWS account can have up to 100 unique event sources of the CloudWatch Events- Schedule source type. Each of these can be the event source for up to five Lambda functions. That is, you can have up to 500 Lambda functions that can be executing on a schedule in your AWS account.

The console also provides a blueprint (lambda-canary) that uses the CloudWatch Events - Schedule source type. Using this blueprint, you can create a sample Lambda function and test this feature. The example code that the blueprint provides checks for the presence of a specific webpage and specific text string on the webpage. If either the webpage or the text string is not found, the Lambda function throws an error.

NEW QUESTION 214
.....

Relate Links

100% Pass Your SOA-C01 Exam with Exam Bible Prep Materials

<https://www.exambible.com/SOA-C01-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>