

SOA-C01 Dumps

AWS Certified SysOps Administrator - Associate

<https://www.certleader.com/SOA-C01-dumps.html>



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

When preparing for a compliance assessment of your system built inside of AWS, what are three best-practices for you to prepare for an audit?

Choose 3 answers

- A. Gather evidence of your IT operational controls
- B. Request and obtain applicable third-party audited AWS compliance reports and certifications
- C. Request and obtain a compliance and security tour of an AWS data center for a pre-assessment security review
- D. Request and obtain approval from AWS to perform relevant network scans and in-depth penetration tests of your system's Instances and endpoints
- E. Schedule meetings with AWS's third-party auditors to provide evidence of AWS compliance that maps to your control objectives

Answer: ABD

NEW QUESTION 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

An organization has configured a VPC with an Internet Gateway (IGW), pairs of public and private subnets (each with one subnet per Availability Zone), and an Elastic Load Balancer (ELB) configured to use the public subnets. The application's web tier leverages the ELB. Auto Scaling and a multi-AZ RDS database instance. The organization would like to eliminate any potential single points of failure in this design. What step should you take to achieve this organization's objective?

- A. Nothing, there are no single points of failure in this architecture.
- B. Create and attach a second IGW to provide redundant internet connectivity.
- C. Create and configure a second Elastic Load Balancer to provide a redundant load balancer.
- D. Create a second multi-AZ RDS instance in another Availability Zone and configure replication to provide a redundant database.

Answer: A

NEW QUESTION 5

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

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

Answer: A

NEW QUESTION 6

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

What option would you implement to successfully launch this application?

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

Answer: C

Explanation:

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

NEW QUESTION 7

An application that you are managing has EC2 instances & DynamoDB tables deployed to several AWS Regions. In order to monitor the performance of the application globally, you would like to see two graphs: 1) Avg CPU Utilization across all EC2 instances and 2) Number of Throttled Requests for all DynamoDB tables.

How can you accomplish this?

- A. Tag your resources with the application name, and select the tag name as the dimension in the CloudWatch Management console to view the respective graphs
- B. Use the Cloud Watch CLI tools to pull the respective metrics from each regional endpoint. Aggregate the data offline & store it for graphing in CloudWatch.
- C. Add SNMP traps to each instance and DynamoDB table. Leverage a central monitoring server to capture data from each instance and table. Put the aggregate data into Cloud Watch for graphing.
- D. Add a CloudWatch agent to each instance and attach one to each DynamoDB table
- E. When configuring the agent, set the appropriate application name & view the graphs in CloudWatch.

Answer: A

Explanation:

Correct answer should be A. When you turn on detailed monitoring in CloudWatch, you can get 1) Avg CPU Utilization across all EC2 instances and 2) Number of Throttled Requests for all DynamoDB tables

Reference: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/GetSingleMetricAllDimensions.html>

NEW QUESTION 8

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

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

Answer: C

Explanation:

But it is possibly A as DynamoDB is automatically available across three facilities in an AWS Region. So moving in to a same AZ is not possible / necessary.

In this case the DB layer is not the issue, the EC2 8xlarge is the issue; so add another one with a ELB in-front of it.
See also: <https://aws.amazon.com/dynamodb/faqs/>

NEW QUESTION 9

You are managing a legacy application inside VPC with hard coded IP addresses in its configuration. Which two mechanisms will allow the application to failover to new instances without the need for reconfiguration? Choose 2 answers

- A. Create an ELB to reroute traffic to a failover instance
- B. Create a secondary ENI that can be moved to a failover instance
- C. Use Route53 health checks to fail traffic over to a failover instance
- D. Assign a secondary private IP address to the primary ENI that can be moved to a failover instance

Answer: BD

Explanation:

This is an odd question. First of all, option A cannot be right because ELB does not failover. Cannot be C because Route 53 does work with hard coded IP. Only B & D cannot be ruled out so best answer.

NEW QUESTION 10

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

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

- A. Denies the server with the IP address 192.168.100.0 full access to the "mybucket" bucket
- B. Denies the server with the IP address 192.168.100.188 full access to the "mybucket" bucket
- C. Grants all the servers within the 192.168.100.0/24 subnet full access to the "mybucket" bucket
- D. Grants all the servers within the 192.168.100.188/32 subnet full access to the "mybucket" bucket

Answer: B

NEW QUESTION 10

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

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 11

When creation of an EBS snapshot is initiated but not completed the EBS volume?

- A. Cannot be detached or attached to an EC2 instance until the snapshot completes
- B. Can be used in read-only mode while the snapshot is in progress
- C. Can be used while the snapshot is in progress
- D. Cannot be used until the snapshot completes

Answer: C

Explanation:

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

NEW QUESTION 16

You are attempting to connect to an instance in Amazon VPC without success. You have already verified that the VPC has an Internet Gateway (IGW) the instance has an associated Elastic IP (EIP) and correct security group rules are in place. Which VPC component should you evaluate next?

- A. The configuration of a NAT instance
- B. The configuration of the Routing Table
- C. The configuration of the internet Gateway (IGW)
- D. The configuration of SRC/DST checking

Answer: B

Explanation:

Reference:

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

NEW QUESTION 20

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

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

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 21

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 22

A user is planning to evaluate AWS for their internal use. The user does not want to incur any charge on his account during the evaluation. Which of the below mentioned AWS services would incur a charge if used?

- A. AWS S3 with 1 GB of storage
- B. AWS micro instance running 24 hours daily
- C. AWS ELB running 24 hours a day
- D. AWS PIOPS volume of 10 GB size

Answer: D

Explanation:

AWS is introducing a free usage tier for one year to help the new AWS customers get started in Cloud. The free tier can be used for anything that the user wants to run in the Cloud. AWS offers a handful of AWS services as a part of this which includes 750 hours of free micro instances and 750 hours of ELB. It includes the AWS S3 of 5 GB and AWS EBS general purpose volume upto 30 GB. PIOPS is not part of free usage tier.

NEW QUESTION 24

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 28

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 32

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

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

Answer: B

Explanation:

The admin can capture information about Elastic Load Balancer using either:

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

NEW QUESTION 35

A user is planning to use AWS CloudFormation. Which of the below mentioned functionalities does not help him to correctly understand CloudFormation?

- A. CloudFormation follows the DevOps model for the creation of Dev & Test
- B. AWS CloudFormation does not charge the user for its service but only charges for the AWS resources created with it.
- C. CloudFormation works with a wide variety of AWS services, such as EC2, EBS, VPC, IAM, S3, RDS, ELB, etc.
- D. CloudFormation provides a set of application bootstrapping scripts which enables the user to install Software.

Answer: A

Explanation:

AWS CloudFormation is an application management tool which provides application modelling, deployment, configuration, management and related activities. It supports a wide variety of AWS services, such as EC2, EBS, AS, ELB, RDS, VPC, etc. It also provides application bootstrapping scripts which enable the user to install software packages or create folders. It is free of the cost and only charges the user for the services created with it. The only challenge is that it does not follow any model, such as DevOps; instead customers can define templates and use them to provision and manage the AWS resources in an orderly way.

NEW QUESTION 40

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 43

A system admin is managing buckets, objects and folders with AWS S3. Which of the below mentioned statements is true and should be taken in consideration by the sysadmin?

- A. The folders support only ACL
- B. Both the object and bucket can have an Access Policy but folder cannot have policy
- C. Folders can have a policy
- D. Both the object and bucket can have ACL but folders cannot have ACL

Answer: A

Explanation:

A sysadmin can grant permission to the S3 objects or the buckets to any user or make objects public using the bucket policy and user policy. Both use the JSON-based access policy language. Generally if user is defining the ACL on the bucket, the objects in the bucket do not inherit it and vice a versa. The bucket policy can be defined at the bucket level which allows the objects as well as the bucket to be public with a single policy applied to that bucket. It cannot be applied at the object level. The folders are similar to objects with no content. Thus, folders can have only ACL and cannot have a policy.

NEW QUESTION 44

An application is generating a log file every 5 minutes. The log file is not critical but may be required only for verification in case of some major issue. The file should be accessible over the internet whenever required. Which of the below mentioned options is a best possible storage solution for it?

- A. AWS S3
- B. AWS Glacier
- C. AWS RDS
- D. AWS RRS

Answer: D

Explanation:

Amazon S3 stores objects according to their storage class. There are three major storage classes: Standard, Reduced Redundancy Storage and Glacier. Standard is for AWS S3 and provides very high durability. However, the costs are a little higher. Glacier is for archival and the files are not available over the internet. Reduced Redundancy Storage is for less critical files. Reduced Redundancy is little cheaper as it provides less durability in comparison to S3. In this case since the log files are not mission critical files, RRS will be a better option.

NEW QUESTION 48

An organization, which has the AWS account ID as 999988887777, has created 50 IAM users. All the users are added to the same group cloudacademy. If the organization has enabled that each IAM user can login with the AWS console, which AWS login URL will the IAM users use?

- A. <https://999988887777.signin.aws.amazon.com/console/>
- B. <https://signin.aws.amazon.com/cloudacademy/>
- C. <https://cloudacademy.signin.aws.amazon.com/999988887777/console/>
- D. <https://999988887777.aws.amazon.com/cloudacademy/>

Answer: A

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Once the organization has created the IAM users, they will have a separate AWS console URL to login to the AWS console. The console login URL for the IAM user will be https://AWS_Account_ID.signin.aws.amazon.com/console/. It uses only the AWS account ID and does not depend on the group or user ID.

NEW QUESTION 50

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 52

A user has configured ELB with three instances. The user wants to achieve High Availability as well as redundancy with ELB. Which of the below mentioned AWS services helps the user achieve this for ELB?

- A. Route 53
- B. AWS Mechanical Turk
- C. Auto Scaling
- D. AWS EMR

Answer: A

Explanation:

The user can provide high availability and redundancy for applications running behind Elastic Load Balancer by enabling the Amazon Route 53 Domain Name System (DNS) failover for the load balancers. Amazon Route 53 is a DNS service that provides reliable routing to the user's infrastructure.

NEW QUESTION 54

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

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

Answer: A

Explanation:

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

NEW QUESTION 57

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 58

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

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

Answer: B

Explanation:

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

NEW QUESTION 59

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

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

Answer: B

Explanation:

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

NEW QUESTION 64

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 65

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

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

Answer: A

Explanation:

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

NEW QUESTION 66

A sys admin is trying to understand EBS snapshots. Which of the below mentioned statements will not be useful to the admin to understand the concepts about a snapshot?

- A. The snapshot is synchronous
- B. It is recommended to stop the instance before taking a snapshot for consistent data
- C. The snapshot is incremental
- D. The snapshot captures the data that has been written to the hard disk when the snapshot command was executed

Answer: A

Explanation:

The AWS snapshot is a point in time backup of an EBS volume. When the snapshot command is executed it will capture the current state of the data that is written on the drive and take a backup. For a better and consistent snapshot of the root EBS volume, AWS recommends stopping the instance. For additional volumes it is recommended to unmount the device. The snapshots are asynchronous and incremental.

NEW QUESTION 71

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

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

Answer: C

Explanation:

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

NEW QUESTION 72

A user has setup an EBS backed instance and a CloudWatch alarm when the CPU utilization is more than 65%. The user has setup the alarm to watch it for 5 periods of 5 minutes each. The CPU utilization is 60% between 9 AM to 6 PM. The user has stopped the EC2 instance for 15 minutes between 11 AM to 11:15 AM. What will be the status of the alarm at 11:30 AM?

- A. Alarm
- B. OK
- C. Insufficient Data
- D. Error

Answer: B

Explanation:

Amazon CloudWatch alarm watches a single metric over a time period the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The state of the alarm will be OK for the whole day. When the user stops the instance for three periods the alarm may not receive the data

NEW QUESTION 77

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

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

Answer: C

Explanation:

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

NEW QUESTION 78

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 81

A user has launched an EC2 instance from an instance store backed AMI. The infrastructure team wants to create an AMI from the running instance. Which of the below mentioned steps will not be performed while creating the AMI?

- A. Define the AMI launch permissions
- B. Upload the bundled volume
- C. Register the AMI
- D. Bundle the volume

Answer: A

Explanation:

When the user has launched an EC2 instance from an instance store backed AMI, it will need to follow certain steps, such as ??Bundling the root volume??. ??Uploading the bundled volume?? and ??Register the AMI??. Once the AMI is created the user can setup the launch permission. However, it is not required to setup during the launch.

NEW QUESTION 84

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

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

Answer: C

Explanation:

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

NEW QUESTION 88

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

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

Answer: A

Explanation:

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

NEW QUESTION 91

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 92

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 96

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

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

Answer: C

Explanation:

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

NEW QUESTION 100

A user has created a VPC with public and private subnets using the VPC wizard. The user has not launched any instance manually and is trying to delete the VPC. What will happen in this scenario?

- A. It will not allow to delete the VPC as it has subnets with route tables
- B. It will not allow to delete the VPC since it has a running route instance
- C. It will terminate the VPC along with all the instances launched by the wizard
- D. It will not allow to delete the VPC since it has a running NAT instance

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. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance with an elastic IP. If the user is trying to delete the VPC it will not allow as the NAT instance is still running.

NEW QUESTION 103

A user has launched an EBS backed instance with EC2-Classic. The user stops and starts the instance. Which of the below mentioned statements is not true with respect to the stop/start action?

- A. The instance gets new private and public IP addresses
- B. The volume is preserved
- C. The Elastic IP remains associated with the instance
- D. The instance may run on a new host computer

Answer: C

Explanation:

A user can always stop/start an EBS backed EC2 instance. When the user stops the instance, it first enters the stopping state, and then the stopped state. AWS does not charge the running cost but charges only for the EBS storage cost. If the instance is running in EC2-Classic, it receives a new private IP address; as the Elastic IP address (EIP) associated with the instance is no longer associated with that instance.

NEW QUESTION 104

An AWS account wants to be part of the consolidated billing of his organization's payee account. How can the owner of that account achieve this?

- A. The payee account has to request AWS support to link the other accounts with his account
- B. The owner of the linked account should add the payee account to his master account list from the billing console
- C. The payee account will send a request to the linked account to be a part of consolidated billing
- D. The owner of the linked account requests the payee account to add his account to consolidated billing

Answer: C

Explanation:

AWS consolidated billing enables the organization to consolidate payments for multiple Amazon Web Services (AWS) accounts within a single organization by making a single paying account. To add a particular account (linked to the master (payee) account, the payee account has to request the linked account to join consolidated billing. Once the linked account accepts the request henceforth all charges incurred by the linked account will be paid by the payee account.

NEW QUESTION 108

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

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

NEW QUESTION 109

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

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

Answer: A

Explanation:

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

NEW QUESTION 112

A user is using the AWS EC2. The user wants to make so that when there is an issue in the EC2 server, such as instance status failed, it should start a new instance in the user's private cloud. Which AWS service helps to achieve this automation?

- A. AWS CloudWatch + Cloudformation
- B. AWS CloudWatch + AWS AutoScaling + AWS ELB
- C. AWS CloudWatch + AWS VPC
- D. AWS CloudWatch + AWS SNS

Answer: D

Explanation:

Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. The user can configure a web service (HTTP End point) in his data centre which receives data and launches an instance in the private cloud. The user should configure the CloudWatch alarm to send a notification to SNS when the `StatusCheckFailed` metric is true for the EC2 instance. The SNS topic can be configured to send a notification to the user's HTTP end point which launches an instance in the private cloud.

NEW QUESTION 116

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

- A. Create an IAM policy with the security group and use that security group for AWS console login
- B. Create an IAM policy with a condition which denies access when the IP address range is not from the organization
- C. Configure the EC2 instance security group which allows traffic only from the organization's IP range
- D. Create an IAM policy with VPC and allow a secure gateway between the organization and AWS Console

Answer: B

Explanation:

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

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

NEW QUESTION 119

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 124

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

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

Answer: A

NEW QUESTION 126

A user has created an Auto Scaling group with default configurations from CLI. The user wants to setup the CloudWatch alarm on the EC2 instances, which are launched by the Auto Scaling group. The user has setup an alarm to monitor the CPU utilization every minute. Which of the below mentioned statements is true?

- A. It will fetch the data at every minute but the four data points [corresponding to 4 minutes] will not have value since the EC2 basic monitoring metrics are collected every five minutes
- B. It will fetch the data at every minute as detailed monitoring on EC2 will be enabled by the default launch configuration of Auto Scaling
- C. The alarm creation will fail since the user has not enabled detailed monitoring on the EC2 instances
- D. The user has to first enable detailed monitoring on the EC2 instances to support alarm monitoring at every minute

Answer: B

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. To enable detailed instance monitoring for a new Auto Scaling group, the user does not need to take any extra steps. When the user creates an Auto Scaling launch config using CLI, each launch configuration contains a flag named InstanceMonitoring.Enabled. The default value of this flag is true. Thus, by default detailed monitoring will be enabled for Auto Scaling as well as for all the instances launched by that Auto Scaling group.

NEW QUESTION 130

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

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

Answer: A

Explanation:

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

NEW QUESTION 134

A user has created a VPC with a subnet and a security group. The user has launched an instance in that subnet and attached a public IP. The user is still unable to connect to the instance. The internet gateway has also been created. What can be the reason for the error?

- A. The internet gateway is not configured with the route table
- B. The private IP is not present
- C. The outbound traffic on the security group is disabled
- D. The internet gateway is not configured with the security group

Answer: A

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. When a user launches an instance and wants to connect to an instance, he needs an internet gateway. The internet gateway should be configured with the route table to allow traffic from the internet.

NEW QUESTION 139

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

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

Answer: D

Explanation:

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

NEW QUESTION 144

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

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

Answer: A

Explanation:

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

NEW QUESTION 149

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

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

Answer: A

Explanation:

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

NEW QUESTION 150

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

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

Answer: B

Explanation:

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

NEW QUESTION 151

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 154

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 156

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

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

Answer: A

Explanation:

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

NEW QUESTION 160

A sys admin is planning to subscribe to the RDS event notifications. For which of the below mentioned source categories the subscription cannot be configured?

- A. DB security group
- B. DB snapshot
- C. DB options group
- D. DB parameter group

Answer: C

Explanation:

Amazon RDS uses the Amazon Simple Notification Service (SNS) to provide a notification when an Amazon RDS event occurs. These events can be configured for source categories, such as DB instance, DB security group, DB snapshot and DB parameter group.

NEW QUESTION 165

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 169

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

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

Answer: D

Explanation:

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

NEW QUESTION 172

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 173

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 174

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

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

Answer: D

Explanation:

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

NEW QUESTION 179

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

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

Answer: B

Explanation:

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

NEW QUESTION 180

A user has configured ELB with SSL using a security policy for secure negotiation between the client and load balancer. The ELB security policy supports various ciphers. Which of the below mentioned options helps identify the matching cipher at the client side to the ELB cipher list when client is requesting ELB DNS over SSL?

- A. Cipher Protocol
- B. Client Configuration Preference
- C. Server Order Preference
- D. Load Balancer Preference

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. When client is requesting ELB DNS over SSL and if the load balancer is configured to support the Server Order Preference, then the load balancer gets to select the first cipher in its list that matches any one of the ciphers in the client's list. Server Order Preference ensures that the load balancer determines which cipher is used for the SSL connection.

NEW QUESTION 183

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

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

Answer: D

Explanation:

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

NEW QUESTION 187

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 190

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

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

Answer: C

Explanation:

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

NEW QUESTION 191

The CFO of a company wants to allow one of his employees to view only the AWS usage report page. Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

- A. "Effect": "Allow", "Action": [??Describe??], "Resource": "Billing"
- B. "Effect": "Allow", "Action": ["AccountUsage"], "Resource": "**"
- C. "Effect": "Allow", "Action": ["aws-portal:ViewUsage"], "Resource": "**"
- D. "Effect": "Allow", "Action": ["aws-portal:ViewBilling"], "Resource": "**"

Answer: C

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow", "Action": [
        "aws-portal:ViewUsage"
      ],
      "Resource": "**"
    }
  ]
}
```

NEW QUESTION 195

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

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

Answer: D

Explanation:

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

```
{
  "Version": "2012-10-17",
```

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

NEW QUESTION 196

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

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

Answer: B

Explanation:

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

NEW QUESTION 198

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

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

Answer: D

NEW QUESTION 200

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 DeletionTerminationProtection 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 204

A user is trying to pre-warm a blank EBS volume attached to a Linux instance. Which of the below mentioned steps should be performed by the user?

- A. There is no need to pre-warm an EBS volume
- B. Contact AWS support to pre-warm
- C. Unmount the volume before pre-warming
- D. Format the device

Answer: C

Explanation:

When the user creates a new EBS volume or restores a volume from the snapshot, the back-end storage blocks are immediately allocated to the user EBS. However, the first time when the user is trying to access a block of the storage, it is recommended to either be wiped from the new volumes or instantiated from the snapshot (for restored volumes) before the user can access the block. This preliminary action takes time and can cause a 5 to 50 percent loss of IOPS for the volume when the block is accessed for the first time. To avoid this it is required to pre-warm the volume. Pre-warming an EBS volume on a Linux instance requires that the user should unmount the blank device first and then write all the blocks on the device using a command, such as `dd`.

NEW QUESTION 205

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

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

Answer: D

Explanation:

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

NEW QUESTION 206

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

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

Answer: A

NEW QUESTION 207

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

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

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

Answer: BC

Explanation:

Reference:

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

NEW QUESTION 209

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 redirected 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 redirected to the appropriate zone to create their profile

Answer: C

NEW QUESTION 211

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

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

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

Answer: D

Explanation:

Reference:

<https://java.awsblog.com/post/Tx1VE22EWF4H86/Accessing-Private-Content-in-Amazon-CloudFront>

NEW QUESTION 215

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

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

Answer: AC

Explanation:

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

NEW QUESTION 219

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

Choose 2 answers

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

Answer: CD

Explanation:

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

NEW QUESTION 223

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 227

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 229

The Security team is concerned because the number of AWS identity and access Management (IAM) policies being in the environment is increasing. The team tasked a SysOps Administrator to report on the number of IAM policies in use and use the total IAM policies.

Which AWS service should the Administrator use to check how current IAM policy compares to current limits?

- A. AWS Trusted Advisor
- B. Amazon Inspector
- C. AWS Config
- D. Organizations

Answer: C

Explanation:

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

NEW QUESTION 231

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 53. 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 for 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 for 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 234

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 237

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

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

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

Answer: C

NEW QUESTION 239

A SysOps Administrator has set up a new Application Load Balancer (ALB) in front of a pair of private web server in multiple Availability Zones. After deployment an updates CloudFormation template with many changes, user now goes to one web server only.

What is the MOST likely reason that the traffic is not being balanced between both servers?

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

Answer: B

Explanation:

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

NEW QUESTION 244

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

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

Answer: D

Explanation:

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

NEW QUESTION 248

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SOA-C01 Exam with Our Prep Materials Via below:

<https://www.certleader.com/SOA-C01-dumps.html>