

Amazon-Web-Services

Exam Questions SOA-C01

AWS Certified SysOps Administrator - Associate



NEW QUESTION 1

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

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

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 2

When attached to an Amazon VPC which two components provide connectivity with external networks? Choose 2 answers

- A. Elastic IPS (EIP)
- B. NAT Gateway (NAT)
- C. Internet Gateway (IGW)
- D. Virtual Private Gateway (VGW)

Answer: CD

NEW QUESTION 3

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 4

You have a web application leveraging an Elastic Load Balancer (ELB) In front of the web servers deployed using an Auto Scaling Group. Your database is running on Relational Database Service (RDS) The application serves out technical articles and responses to them in general there are more views of an article than there are responses to the article. On occasion, an article on the site becomes extremely popular resulting in significant traffic Increases that causes the site to go down.

What could you do to help alleviate the pressure on the infrastructure while maintaining availability during these events?

Choose 3 answers

- A. Leverage CloudFront for the delivery of the articles.
- B. Add RDS read-replicas for the read traffic going to your relational database
- C. Leverage ElastiCache for caching the most frequently used data.
- D. Use SQS to queue up the requests for the technical posts and deliver them out of the queue.
- E. Use Route53 health checks to fail over to an S3 bucket for an error page.

Answer: ABC

Explanation:

The questions mention RDS so an answer that includes that as part of the solution makes sense. Also, Route53 does nothing to alleviate pressure on the infrastructure, it's for failover. E is counterproductive. It talks about failing over to an error page on S3.

NEW QUESTION 5

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 LOAP 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 LOAP 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 6

You have a Linux EC2 web server instance running inside a VPC. The instance is in a public subnet and has an EIP associated with it so you can connect to it over the Internet via HTTP or SSH. The instance was also fully accessible when you last logged in via SSH, and was also serving web requests on port 80. Now you are not able to SSH into the host nor does it respond to web requests on port 80 that were working fine last time you checked. You have double-checked that all networking configuration parameters (security groups, route tables, IGW, EIP, NACLs, etc) are properly configured (and you haven't made any changes to those anyway since you were last able to reach the instance). You look at the EC2 console and notice that system status check shows "impaired." Which should be your next step in troubleshooting and attempting to get the instance back to a healthy state so that you can log in again?

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

Answer: A

NEW QUESTION 7

You have two Elastic Compute Cloud (EC2) instances inside a Virtual Private Cloud (VPC) in the same Availability Zone (AZ) but in different subnets. One instance is running a database and the other instance an application that will interface with the database. You want to confirm that they can talk to each other for your application to work properly.

Which two things do we need to confirm in the VPC settings so that these EC2 instances can communicate inside the VPC?

Choose 2 answers

- A. A network ACL that allows communication between the two subnets.
- B. Both instances are the same instance class and using the same Key-pair.
- C. That the default route is set to a NAT instance or internet Gateway (IGW) for them to communicate.
- D. Security groups are set to allow the application host to talk to the database on the right port/protocol.

Answer: AD

NEW QUESTION 8

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

Choose 2 answers

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

Answer: AC

Explanation:

The below services provide Root level access:

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

NEW QUESTION 9

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 10

You run a web application where web servers on EC2 instances are in an Auto Scaling group. Monitoring over the last 6 months shows that 6 web servers are necessary to handle the minimum load. During the day up to 12 servers are needed. Five to six days per year, the number of web servers required might go up to 15.

What would you recommend to minimize costs while being able to provide high availability?

- A. 6 Reserved instances (heavy utilization), 6 Reserved instances (medium utilization), rest covered by On-Demand instances
- B. 6 Reserved instances (heavy utilization), 6 On-Demand instances, rest covered by Spot Instances
- C. 6 Reserved instances (heavy utilization), 6 Spot instances, rest covered by On-Demand instances
- D. 6 Reserved instances (heavy utilization), 6 Reserved instances (medium utilization), rest covered by Spot instances

Answer: A

Explanation:

The only plausible answer is A because all other answers include Spot Instances that can be removed without warning and that would not be highly available.

NEW QUESTION 10

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

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

Answer: B

NEW QUESTION 12

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 15

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

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

Answer: C

Explanation:

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

NEW QUESTION 19

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 22

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 25

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

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

Answer: A

Explanation:

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

NEW QUESTION 29

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

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

Answer: B

Explanation:

AWS consolidated billing enables the organization to consolidate payments for multiple Amazon Web Services (AWS) accounts within a single organization by making a single paying account. Consolidated billing enables the organization to see a combined view of the AWS charges incurred by each account as well as obtain a detailed cost report for each of the individual AWS accounts associated with the paying account. The payee account will not have any other access than billing data of linked accounts.

NEW QUESTION 32

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 37

An organization is planning to use AWS for their production roll out. The organization wants to implement automation for deployment such that it will automatically create a LAMP stack, download the latest PHP installable from S3 and setup the ELB. Which of the below mentioned AWS services meets the requirement for making an orderly deployment of the software?

- A. AWS Elastic Beanstalk
- B. AWS CloudFront
- C. AWS CloudFormation
- D. AWS DevOps

Answer: C

Explanation:

AWS CloudFormation is an application management tool which provides application modelling, deployment, configuration, management and related activities. CloudFormation provides an easy way to create and delete the collection of related AWS resources and provision them in an orderly way. AWS CloudFormation automates and simplifies the task of repeatedly and predictably creating groups of related resources that power the user's applications. AWS CloudFront is a CDN; Elastic Beanstalk does quite a few of the required tasks. However, it is a PaaS which uses a ready AMI. AWS Elastic Beanstalk provides an environment to easily develop and run applications in the cloud.

NEW QUESTION 41

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

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

Answer: B

Explanation:

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

NEW QUESTION 46

A user has configured the Auto Scaling group with the minimum capacity as 3 and the maximum capacity as 5. When the user configures the AS group, how many instances will Auto Scaling launch?

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

Answer: A

Explanation:

The default launch is going to configure 3 instances. It will only go to 5 under certain conditions. It won't launch initially 5 instances. This question doesn't give enough information to answer fully.

NEW QUESTION 51

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

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

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html CloudWatch monitors the following services. As soon as you begin using a service, it automatically sends metrics to CloudWatch for you.

CloudWatch offers either basic or detailed monitoring for supported AWS products. Basic monitoring means that a service sends data points to CloudWatch every five minutes. Detailed monitoring means that a service sends data points to CloudWatch every minute.

Note

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

Auto Scaling

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

NEW QUESTION 52

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

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

Answer: B

Explanation:

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

NEW QUESTION 55

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 58

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

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

Answer: B

Explanation:

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

NEW QUESTION 62

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

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

Answer: D

Explanation:

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

NEW QUESTION 65

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 69

A user has created numerous EBS volumes. What is the general limit for each AWS account for the maximum number of EBS volumes that can be created?

- A. 10000
- B. 5000
- C. 100
- D. 1000

Answer: B

Explanation:

A user can attach multiple EBS volumes to the same instance within the limits specified by his AWS account. Each AWS account has a limit on the number of Amazon EBS volumes that the user can create, and the total storage available. The default limit for the maximum number of volumes that can be created is 5000.

NEW QUESTION 73

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

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

Answer: C

Explanation:

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy, eliminate I/O freezes, and minimize latency spikes during system backups. Running a DB instance with high availability can enhance availability during planned system maintenance, and help protect your databases against DB instance failure and Availability Zone disruption. Note that the high-availability feature is not a scaling solution for read-only scenarios; you cannot use a standby replica to serve read traffic. To service read-only traffic, you should use a read replica.

NEW QUESTION 74

An organization is using cost allocation tags to find the cost distribution of different departments and projects. One of the instances has two separate tags with the key/ value as ??InstanceName/HR??. ??CostCenter/HR??. What will AWS do in this case?

- A. InstanceName is a reserved tag for AW
- B. Thus, AWS will not allow this tag
- C. AWS will not allow the tags as the value is the same for different keys
- D. AWS will allow tags but will not show correctly in the cost allocation report due to the same value of the two separate keys
- E. AWS will allow both the tags and show properly in the cost distribution report

Answer: D

Explanation:

AWS provides cost allocation tags to categorize and track the AWS costs. When the user applies tags to his AWS resources, AWS generates a cost allocation report as a comma-separated value (CSV file) with the usage and costs aggregated by those tags. Each tag will have a key-value and can be applied to services, such as EC2, S3, RDS, EMR, etc. It is required that the key should be different for each tag. The value can be the same for different keys. In this case since the value is different, AWS will properly show the distribution report with the correct values.

NEW QUESTION 75

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

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

Answer: B

Explanation:

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

NEW QUESTION 80

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 82

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 87

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

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

Answer: C

Explanation:

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

NEW QUESTION 92

An AWS root account owner is trying to create a policy to access RDS. Which of the below mentioned statements is true with respect to the above information?

- A. Create a policy which allows the users to access RDS and apply it to the RDS instances
- B. The user cannot access the RDS database if he is not assigned the correct IAM policy
- C. The root account owner should create a policy for the IAM user and give him access to the RDS services
- D. The policy should be created for the user and provide access for RDS

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 account owner wants to create a policy for RDS, the owner has to create an IAM user and define the policy which entitles the IAM user with various RDS services such as Launch Instance, Manage security group, Manage parameter group etc.

NEW QUESTION 95

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

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

Answer: C

Explanation:

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

An Availability Zone outage
The primary DB instance fails

The DB instance's server type is changed

The DB instance is undergoing software patching

A manual failover of the DB instance was initiated using Reboot with failover

NEW QUESTION 99

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

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

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

Answer: D

Explanation:

A sysadmin can grant permission to the S3 objects or the buckets to any user or make objects public using the bucket policy and user policy. Both use the JSON-based access policy language. Generally if the user is defining the ACL on the bucket, the objects in the bucket do not inherit it and vice versa. The bucket policy can be defined at the bucket level which allows the objects as well as the bucket to be public with a single policy applied to that bucket. In the sample policy the action says "s3:ListBucket" for effect Allow on Resource arn:aws:s3:::cloudacademy. This will make the cloudacademy bucket public.

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

NEW QUESTION 104

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

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

Answer: C

Explanation:

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

NEW QUESTION 106

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 110

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

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

Answer: D

Explanation:

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

NEW QUESTION 113

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 117

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. The NAT instance ID is i-a12345. Which of the below mentioned entries are required in the main route table attached with the private subnet to allow instances to connect with the internet?

- A. Destination: 0.0.0.0/0 and Target: i-a12345
- B. Destination: 20.0.0.0/0 and Target: 80
- C. Destination: 20.0.0.0/0 and Target: i-a12345
- D. Destination: 20.0.0.0/24 and Target: i-a12345

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 two route tables and attach to the subnets. The main route table will have the entry ??Destination: 0.0.0.0/0 and Target: ia12345??. which allows all the instances in the private subnet to connect to the internet using NAT.

NEW QUESTION 122

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 126

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

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

Answer: A

Explanation:

CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. Services, such as RDS, EC2, Auto Scaling, ELB, and Route 53 can provide the monitoring data every minute.

NEW QUESTION 128

A user is measuring the CPU utilization of a private data centre machine every minute. The machine provides the aggregate of data every hour, such as Sum of data??. ??Min value??. ??Max value, and ??Number of Data points??.

The user wants to send these values to CloudWatch. How can the user achieve this?

- A. Send the data using the put-metric-data command with the aggregate-values parameter
- B. Send the data using the put-metric-data command with the average-values parameter
- C. Send the data using the put-metric-data command with the statistic-values parameter
- D. Send the data using the put-metric-data command with the aggregate ?Vdata parameter

Answer: C

Explanation:

AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. The user can publish the data to CloudWatch as single data points or as an aggregated set of data points called a statistic set using the command put- metric-data. When sending the aggregate data, the user needs to send it with the parameter statistic-values:

```
awscloudwatch put-metric-data --metric-name <Name> --namespace <Custom namespace -- timestamp  
<UTC Format> --statistic-values Sum=XX,Minimum=YY,Maximum=AA,SampleCount=BB --unit Milliseconds
```

NEW QUESTION 133

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

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

Answer: C

Explanation:

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

NEW QUESTION 135

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

- A. Which DB is the user using right now?
- B. My SQL

- C. Oracle
- D. MS SQL
- E. PostgreSQL

Answer: C

Explanation:

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

NEW QUESTION 139

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

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

Answer: D

Explanation:

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

NEW QUESTION 143

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

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

Answer: C

Explanation:

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

NEW QUESTION 147

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 149

A user has a weighing plant. The user measures the weight of some goods every 5 minutes and sends data to AWS CloudWatch for monitoring and tracking. Which of the below mentioned parameters is mandatory for the user to include in the request list?

- A. Value
- B. Namespace
- C. Metric Name
- D. Timezone

Answer: B

Explanation:

AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. The user can publish the data to CloudWatch as single data points or as an aggregated set of data points called a statistic set. The user has to always include the namespace as part of the request. The user can supply a file instead of the metric name. If the user does not supply the timezone, it accepts the current time. If the user is sending the data as a single data point it will have parameters, such as value. However, if the user is sending as an aggregate it will have parameters, such as statistic-values.

NEW QUESTION 153

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

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

Answer: D

Explanation:

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

NEW QUESTION 157

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 158

A user has launched an EBS backed EC2 instance in the US-East-1a region. The user stopped the instance and started it back after 20 days. AWS throws up an `InsufficientInstanceCapacity` error. What can be the possible reason for this?

- A. AWS does not have sufficient capacity in that availability zone
- B. AWS zone mapping is changed for that user account
- C. There is some issue with the host capacity on which the instance is launched
- D. The user account has reached the maximum EC2 instance limit

Answer: A

Explanation:

When the user gets an `InsufficientInstanceCapacity` error while launching or starting an EC2 instance, it means that AWS does not currently have enough available capacity to service the user request. If the user is requesting a large number of instances, there might not be enough server capacity to host them. The user can either try again later, by specifying a smaller number of instances or changing the availability zone if launching a fresh instance.

NEW QUESTION 160

A system admin wants to add more zones to the existing ELB. The system admin wants to perform this activity from CLI. Which of the below mentioned command helps the system admin to add new zones to the existing ELB?

- A. `elb-enable-zones-for-lb`
- B. `elb-add-zones-for-lb`
- C. It is not possible to add more zones to the existing ELB
- D. `elb-configure-zones-for-lb`

Answer: A

Explanation:

The user has created an Elastic Load Balancer with the availability zone and wants to add more zones to the existing ELB. The user can do so in two ways: From the console or CLI, add new zones to ELB;

NEW QUESTION 163

A user is planning to scale up an application by 8 AM and scale down by 7 PM daily using Auto Scaling. What should the user do in this case?

- A. Setup the scaling policy to scale up and down based on the CloudWatch alarms
- B. The user should increase the desired capacity at 8 AM and decrease it by 7 PM manually
- C. The user should setup a batch process which launches the EC2 instance at a specific time
- D. Setup scheduled actions to scale up or down at a specific time

Answer: A

Explanation:

Auto Scaling based on a schedule allows the user to scale the application in response to predictable load changes. To configure the Auto Scaling group to scale based on a schedule, the user needs to create scheduled actions. A scheduled action tells Auto Scaling to perform a scaling action at a certain time in the future.

NEW QUESTION 166

A sys admin has enabled a log on ELB. Which of the below mentioned activities are not captured by the log?

- A. Response processing time
- B. Front end processing time
- C. Backend processing time
- D. Request processing time

Answer: B

Explanation:

Elastic Load Balancing access logs capture detailed information for all the requests made to the load balancer. Each request will have details, such as client IP, request path, ELB IP, time, and latencies. The time will have information, such as Request Processing time, Backend Processing time and Response Processing time.

NEW QUESTION 167

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

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

Answer: D

Explanation:

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

NEW QUESTION 170

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

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

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

Answer: B

NEW QUESTION 175

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

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

Answer: A

Explanation:

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

NEW QUESTION 178

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

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

Answer: A

Explanation:

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

NEW QUESTION 181

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 182

Your business is building a new application that will store its entire customer database on a RDS MySQL database, and will have various applications and users that will query that data for different purposes.

Large analytics jobs on the database are likely to cause other applications to not be able to get the query results they need to, before time out. Also, as your data grows, these analytics jobs will start to take more time, increasing the negative effect on the other applications.

How do you solve the contention issues between these different workloads on the same data?

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

Answer: B

Explanation:

Amazon ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud. Amazon ElastiCache improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. The service simplifies and offloads the management, monitoring and operation of in-memory cache environments, enabling your engineering resources to focus on developing applications. Using Amazon ElastiCache, you can not only improve load and response times to user actions and queries, but also reduce the cost associated with scaling web applications.

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

NEW QUESTION 187

How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

- A. Simply create a new volume in the other AZ and specify the original volume as the source.
- B. Detach the volume, then use the `ec2-migrate-volume` command to move it to another AZ.
- C. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
- D. Detach the volume and attach it to another EC2 instance in the other AZ.

Answer: C

Explanation:

Snapshots can be used to create multiple new EBS volumes, expand the size of a volume, or move volumes across Availability Zones.

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

NEW QUESTION 190

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

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 194

An application you maintain consists of multiple EC2 instances in a default tenancy VPC. This application has undergone an internal audit and has been determined to require dedicated hardware for one instance. Your compliance team has given you a week to move this instance to single-tenant hardware. Which process will have minimal impact on your application while complying with this requirement?

- A. Create a new VPC with `tenancy=dedicated` and migrate to the new VPC
- B. Use `ec2-reboot-instances` command line and set the parameter `"dedicated=true"`

- C. Right click on the instance, select properties and check the box for dedicated tenancy
- D. Stop the instance, create an AMI, launch a new instance with tenancy=dedicated, and terminate the old instance

Answer: D

Explanation:

Reference:

See: <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html#dedicated-apichanges>

You cannot change the tenancy of a default instance after you've launched it. You can change the tenancy of an instance from ??dedicated?? to ??host?? after you've launched it, and vice versa.

NEW QUESTION 195

A .NET application that you manage is running in Elastic Beanstalk. Your developers tell you they will need access to application log files to debug issues that arise. The infrastructure will scale up and down.

How can you ensure the developers will be able to access only the log files?

- A. Access the log files directly from Elastic Beanstalk
- B. Enable log file rotation to S3 within the Elastic Beanstalk configuration
- C. Ask your developers to enable log file rotation in the applications web.config file
- D. Connect to each Instance launched by Elastic Beanstalk and create a Windows Scheduled task to rotate the log files to S3.

Answer: A

Explanation:

Reference:

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.loggingS3.title.html>

NEW QUESTION 196

A company needs to monitor the read and write IOPs metrics for their AWS MySQL RDS instance and send real-time alerts to their operations team. Which AWS services can accomplish this? Choose 2 answers

- A. Amazon Simple Email Service
- B. Amazon CloudWatch
- C. Amazon Simple Queue Service
- D. Amazon Route 53
- E. Amazon Simple Notification Service

Answer: BE

NEW QUESTION 198

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 200

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 201

A syslog Administrator is created additional Amazon EC2 instances and receive an Instancelimitexceeded error.

What is the cause of the issue and how can it be resolve?

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 205

Based on the AWS Shared Responsibility Model, which of the following actions are the responsibility of the customer for an Aurora database?

- A. Performing underlying OS updates
- B. Provisioning of storage for database
- C. Scheduling maintenance, patches and other updates
- D. Executing maintenance, patches and other updates

Answer: B

NEW QUESTION 209

A company Development team to access the AWS Management Console. A System Administrator has been asked to find a solution so that the Developers can sign in to the console using Active Directory (AD) credentials and not as IAM users.

What steps should the Systems Administrator take to enable functionality?

- A. Set up an Amazon Cognito federation, and then obtain temporary credentials using AWS Security Token Service
- B. Assign the temporary credentials to an IAM role to allow a developer's access to the AWS resource.
- C. Set up Active Directory Connector to use the corporate AD servers. Enable AWS console access under the AWS Directory Service Console for the AD Connector that was just created
- D. Create a role with the resources and permissions that the Development team should have access to use.
- E. Connect the corporate AD servers to AWS using Amazon Cognito user pools. Enable AWS console access within Cognito, and then assign the appropriate role to the user pool.
- F. Create a SAML template file using IAM. Assign the template to the corporate AD through the Simple AD. Grant the Development team access to the SAML template.

Answer: A

NEW QUESTION 212

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 be 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 213

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

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

Answer: D

Explanation:

Getting Credential Reports for Your AWS Account

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

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

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

NEW QUESTION 218

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

- A. AWS CloudTrail
- B. Amazon Inspector
- C. AWS Config
- D. AWS Systems Manager

Answer: C

Explanation:

AWS Config

TRACK RESOURCE INVENTORY AND CHANGES

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

NEW QUESTION 221

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

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

Answer: D

Explanation:

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

NEW QUESTION 223

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 228

A SysOps Administrator has been tasked with deploying a company infrastructure as code. The administrator wants to write a single template that can be reused for multiple environment in a safe, repeatable manner.

What is the recommended way to use AWS CloudFormation to meet this requirement?

- A. Use parameters to provision the resource.
- B. Use nested stack to provision the resources.
- C. Use Amazon EC2 user data to provision the resources.
- D. Use stack policies to provision the resources.

Answer: D

NEW QUESTION 229

A company operate a secure website running an Amazon EC2 instance behind a Classic Load Balancer. An SSL certificate from AWS Certificate Manager is deployment on the load balancer. The company's Marketing team has determined that too many customer using older browser are experiencing issues with the website has asked a SysOps Administrator to fix this issue.

What course of action should the administrator take?

- A. Update the SSL negotiation configuration of the load balancer by creating a custom security policy
- B. Ensure the appropriate cipher has been enabled so that the web application can support the web browser.
- C. Create a separate Classic Load Balancer and install custom SSL certificate with a different domain name on it that support the web browser
- D. Ask customer with the affected browser to use this domain name instead of the one they are accustomed to using.
- E. Create a new SSL certificate in Certificate Manager and install this certificate on each of the servers to accommodate the web browsers.
- F. Remove the load balancer from the configuration and instead install a custom SSL certificate on each of the web servers.

Answer: A

Explanation:

Update the SSL Negotiation Configuration of Your Classic Load Balancer

Elastic Load Balancing provides security policies that have predefined SSL negotiation configurations to use to negotiate SSL connections between clients and your load balancer. If you are using the HTTPS/SSL protocol for your listener, you can use one of the predefined security policies, or use your own custom security policy.

For more information about the security policies, see [SSL Negotiation Configurations for Classic Load Balancers](#). For information about the configurations of the security policies provided by Elastic Load Balancing, see [Predefined SSL Security Policies](#).

If you create an HTTPS/SSL listener without associating a security policy, Elastic Load Balancing associates the default predefined security policy, `ELBSecurityPolicy-2016-08`, with your load balancer. If you have an existing load balancer with an SSL negotiation configuration that does not use the latest protocols and ciphers, we recommend that you update your load balancer to use `ELBSecurityPolicy-2016-08`. If you prefer, you can create a custom configuration. We strongly recommend that you test the new security policies before you upgrade your load balancer configuration.

The following examples show you how to update the SSL negotiation configuration for an HTTPS/SSL listener. Note that the change does not affect requests that were received by a load balancer node and are pending routing to a healthy instance, but the updated configuration will be used with new requests that are received.

NEW QUESTION 232

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SOA-C01 Practice Exam Features:

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

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