

Exam question from Amazon's AWS-SysOps

Question #: 1

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 2

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 three.)

- 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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 3

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You have started a new job and are reviewing your company's infrastructure on AWS. You notice one web application where they have an Elastic Load Balancer (&B) in front of web instances in an Auto Scaling Group. When you check the metrics for the ELB in CloudWatch, you see four healthy instances in Availability Zone (AZ) A and zero in AZ B. There are zero unhealthy instances.

What do you need to fix to balance the instances across AZs?

- A. Set the ELB to only be attached to another AZ
- B. Make sure Auto Scaling is configured to launch in both AZs
- C. Make sure your AMI is available in both AZs
- D. Make sure the maximum size of the Auto Scaling Group is greater than 4

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 4

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You have been asked to leverage Amazon VPC BC2 and SOS to implement an application that submits and receives millions of messages per second to a message queue. You want to ensure your application has sufficient bandwidth between your EC2 instances and SQS.

Which option will provide the most scalable solution for communicating between the application and SQS?

- A. Ensure the application instances are properly configured with an Elastic Load Balancer
- B. Ensure the application instances are launched in private subnets with the EBS-optimized option enabled
- C. Ensure the application instances are launched in public subnets with the associate-public-IP-address=true option enabled
- D. Launch application instances in private subnets with an Auto Scaling group and Auto Scaling triggers configured to watch the SQS queue size

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 5

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 6

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

When attached to an Amazon VPC, which two components provide connectivity with external networks? (Choose two.)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 7

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 8

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You have an Auto Scaling group associated with an Elastic Load Balancer (ELB). You have noticed that instances launched via the Auto Scaling group are being marked unhealthy due to an ELB health check, but these unhealthy instances are not being terminated.

What do you need to do to ensure that instances marked unhealthy by the ELB will be terminated and replaced?

- A. Change the thresholds set on the Auto Scaling group health check
- B. Add an Elastic Load Balancing health check to your Auto Scaling group
- C. Increase the value for the Health check interval set on the Elastic Load Balancer
- D. Change the health check set on the Elastic Load Balancer to use TCP rather than HTTP checks

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 9

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

- A. Amazon S3
- B. Amazon RDS
- C. Amazon EBS
- D. Amazon Red shift

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 10

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 11

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following are characteristics of Amazon VPC subnets? (Choose two.)

- A. Each subnet maps to a single Availability Zone
- B. A CIDR block mask of /25 is the smallest range supported
- C. Instances in a private subnet can communicate with the internet only if they have an Elastic IP.
- D. By default, all subnets can route between each other, whether they are private or public
- E. V Each subnet spans at least 2 Availability zones to provide a high-availability environment

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 12

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 PutMetricData permission and modify the Auto Scaling launch configuration to inject the userscredentials 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 PutMetricData permission and put the credentials in a private repository and have applications on the server pull the credentials as needed

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 13

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 14

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 three.)

- 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 SOS 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.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 15

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 application1?

- 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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 16

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You need to design a VPC for a web-application consisting of an Elastic Load Balancer (ELB), a fleet of web/application servers, and an RDS database. The entire Infrastructure must be distributed over 2 availability zones.

Which VPC configuration works while assuring the database is not available from the Internet?

- A. One public subnet for ELB one public subnet for the web-servers, and one private subnet for the database
- B. One public subnet for ELB two private subnets for the web-servers, two private subnets for RDS
- C. Two public subnets for ELB two private subnets for the web-servers and two private subnets for RDS
- D. Two public subnets for ELB two public subnets for the web-servers, and two public subnets for RDS

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 17

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An application that you are managing has EC2 instances & Dynamo DB 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
- 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. When configuring the agent set the appropriate application name & view the graphs in CloudWatch.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 18

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

When assessing an organization's use of AWS API access credentials which of the following three credentials should be evaluated? (Choose three.)

- A. Key pairs
- B. Console passwords
- C. Access keys
- D. Signing certificates
- E. Security Group memberships

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 19

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 20

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is a placement group?

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 21

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

- A. No, two instances in two different AZ's can't talk directly to each other via ICMP ping as that protocol is not allowed across subnet (iebroadcast) boundaries
- B. Yes, both the monitoring instance and the application instance have to be a part of the same security group, and that security group needs to allow inbound ICMP
- C. Yes, the security group for the monitoring instance needs to allow outbound ICMP and the application instance's security group needs to allow Inbound ICMP
- D. Yes, both the monitoring instance's security group and the application instance's security group need to allow both inbound and outbound ICMP ping packets since ICMP is not a connection-oriented protocol

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 22

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 two.)

- 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.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 23

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which services allow the customer to retain full administrative privileges of the underlying EC2 instances? (Choose two.)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 24

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 25

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 two.)

- 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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 26

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You are designing a system that has a Bastion host. This component needs to be highly available without human intervention.

Which of the following approaches would you select?

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 27

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 28

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 29

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 30

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 31

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You have set up Individual AWS accounts for each project. You have been asked to make sure your AWS Infrastructure costs do not exceed the budget set per project for each month.

Which of the following approaches can help ensure that you do not exceed the budget each month?

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 32

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 33

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You are using ElastiCache Memcached to store session state and cache database queries in your infrastructure. You notice in CloudWatch that Evictions and GetMisses are both very high.

What two actions could you take to rectify this? (Choose two.)

- A. Increase the number of nodes in your cluster
- B. Tweak the max_item_size parameter
- C. Shrink the number of nodes in your cluster
- D. Increase the size of the nodes in the cluster

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 34

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You are running a database on an EC2 instance, with the data stored on Elastic Block Store (EBS) for persistence. At times throughout the day, you are seeing large variance in the response times of the database queries. Looking into the instance with the `isolate` command you see a lot of wait time on the disk volume that the database's data is stored on.

What two ways can you improve the performance of the database's storage while maintaining the current persistence of the data? (Choose two.)

- A. Move to an SSD backed instance
- B. Move the database to an EBS-Optimized Instance
- C. Use Provisioned IOPS EBS
- D. Use the ephemeral storage on an m2 4xlarge Instance Instead

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 35

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 36

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You have decided to change the Instance type for instances running in your application tier that are using Auto Scaling.

In which area below would you change the instance type definition?

- A. Auto Scaling launch configuration
- B. Auto Scaling group
- C. Auto Scaling policy
- D. Auto Scaling tags

[Show Suggested Answer](#)



Exam question from Amazon's AWS-SysOps

Question #: 37

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 38

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Which deployment option meets these requirements while minimizing administrative burden?

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 39

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You have been asked to automate many routine systems administrator backup and recovery activities. Your current plan is to leverage AWS-managed solutions as much as possible and automate the rest with the AWS CLI and scripts.

Which task would be best accomplished with a script?

- A. Creating daily EBS snapshots with a monthly rotation of snapshots
- B. Creating daily RDS snapshots with a monthly rotation of snapshots
- C. Automatically detect and stop unused or underutilized EC2 instances
- D. Automatically add Auto Scaled EC2 instances to an Amazon Elastic Load Balancer

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 40

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Your organization's security policy requires that all privileged users either use frequently rotated passwords or one-time access credentials in addition to username/password.

Which two of the following options would allow an organization to enforce this policy for AWS users? (Choose two.)

- A. Configure multi-factor authentication for privileged IAM users
- B. Create IAM users for privileged accounts
- C. Implement identity federation between your organization's Identity provider leveraging the IAM Security Token Service
- D. Enable the IAM single-use password policy option for privileged users

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 41

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What are characteristics of Amazon S3? (Choose two.)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 42

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You receive a frantic call from a new DBA who accidentally dropped a table containing all your customers.

Which Amazon RDS feature will allow you to reliably restore your database to within 5 minutes of when the mistake was made?

- A. Multi-AZ RDS
- B. RDS snapshots
- C. RDS read replicas
- D. RDS automated backup

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 43

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 44

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 ELB. RDS and EC2 instances

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 45

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You use S3 to store critical data for your company. Several users within your group currently have full permissions to your S3 buckets. You need to come up with a solution that does not impact your users and also protect against the accidental deletion of objects.

Which two options will address this issue? (Choose two.)

- A. Enable versioning on your S3 Buckets
- B. Configure your S3 Buckets with MFA delete
- C. Create a Bucket policy and only allow read only permissions to all users at the bucket level
- D. Enable object life cycle policies and configure the data older than 3 months to be archived in Glacier

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 46

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 47

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You are tasked with setting up a cluster of EC2 Instances for a NoSQL database. The database requires random read I/O 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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 48

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 49

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Which approach addresses this requirement?

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 50

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 51

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Your company is moving towards tracking web page users with a small tracking

image loaded on each page. Currently you are serving this image out of US-East, but are starting to get concerned about the time it takes to load the image for users on the west coast.

What are the two best ways to speed up serving this image? (Choose two.)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 52

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

If you want to launch Amazon Elastic Compute Cloud (EC2) Instances and assign each Instance a predetermined private IP address you should:

- A. Assign a group or sequential Elastic IP address to the instances
- B. Launch the instances in a Placement Group
- C. Launch the instances in the Amazon virtual Private Cloud (VPC).
- D. Use standard EC2 instances since each instance gets a private Domain Name Service (DNS) already
- E. Launch the Instance from a private Amazon Machine image (Mil)

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 53

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A customer has a web application that uses cookie Based sessions to track logged in users. It is deployed on AWS using ELB and Auto Scaling. The customer observes that when load increases. Auto Scaling launches new Instances but the load on the existing Instances does not decrease, causing all existing users have a sluggish experience.

Which two answer choices independently describe a behavior that could be the cause of the sluggish user experience? (Choose two.)

- A. ELB's normal behavior sends requests from the same user to the same backend instance
- B. ELB's behavior when sticky sessions are enabled causes ELB to send requests in the same session to the same backend instance
- C. A faulty browser is not honoring the TTL of the ELB DNS name
- D. The web application uses long polling such as comet or websockets. Thereby keeping a connection open to a web server for a long time

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 54

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

How can the domain's zone apex for example "myzoneapexdomain com" be pointed towards an Elastic Load Balancer?

- A. By using an AAAA record
- B. By using an A record
- C. By using an Amazon Route 53 CNAME record
- D. By using an Amazon Route 53 Alias record

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 55

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 groups. The user can login with their own group ID

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 56

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 57

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 58

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 59

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is planning to use AWS Cloud formation for his automatic deployment requirements. Which of the below mentioned components are required as a part of the template?

- A. Parameters
- B. Outputs
- C. Template version
- D. Resources

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 60

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 61

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has launched an EC2 instance. The user is planning to setup the CloudWatch alarm. Which of the below mentioned actions is not supported by the CloudWatch alarm?

- A. Notify the Auto Scaling launch config to scale up
- B. Send an SMS using SNS
- C. Notify the Auto Scaling group to scale down
- D. Stop the EC2 instance

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 62

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 63

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 64

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has deployed an application on his private cloud. The user is using his own monitoring tool. He wants to configure that whenever there is an error, the monitoring tool should notify him via SMS. Which of the below mentioned AWS services will help in this scenario?

- A. None because the user infrastructure is in the private cloud/
- B. AWS SNS
- C. AWS SES
- D. AWS SMS

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 65

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 an instance by 8 AM and remove it by Friday 6 PM

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 66

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 67

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 68

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has setup a billing alarm using CloudWatch for \$200. The usage of AWS exceeded \$200 after some days. The user wants to increase the limit from \$200 to \$400? What should the user do?

- A. Create a new alarm of \$400 and link it with the first alarm
- B. It is not possible to modify the alarm once it has crossed the usage limit
- C. Update the alarm to set the limit at \$400 instead of \$200
- D. Create a new alarm for the additional \$200 amount

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 69

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

[1]

does this policy define?

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

- A. It is not possible to define a policy at the object level
- B. It will make all the objects of the bucket cloudacademy as public
- C. It will make the bucket cloudacademy as public [1]

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 70

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 71

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 72

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 73

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 74

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 75

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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. 0
- C. 5
- D. 2

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 76

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 77

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 78

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 79

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 80

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 81

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 82

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 83

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 84

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 85

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 86

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 87

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created an ELB with three instances. How many security groups will ELB create by default?

- A. 3
- B. 5
- C. 2
- D. 1

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 88

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 89

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly. Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

- A. AWS Glacier
- B. AWS Elastic Transcoder
- C. AWS Simple Notification Service
- D. AWS Simple Queue Service

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 90

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 91

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 92

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 93

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A sys admin is maintaining an application on AWS. The application is installed on EC2 and user has configured ELB and Auto Scaling. Considering future load increase, the user is planning to launch new servers proactively so that they get registered with ELB. How can the user add these instances with Auto Scaling?

- A. Increase the desired capacity of the Auto Scaling group
- B. Increase the maximum limit of the Auto Scaling group
- C. Launch an instance manually and register it with ELB on the fly
- D. Decrease the minimum limit of the Auto Scaling group

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 94

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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/`

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 95

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has setup connection draining with ELB to allow in-flight requests to continue while the instance is being deregistered through Auto Scaling. If the user has not specified the draining time, how long will ELB allow inflight requests traffic to continue?

- A. 600 seconds
- B. 3600 seconds
- C. 300 seconds
- D. 0 seconds

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 96

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 97

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A sys admin has created a shopping cart application and hosted it on EC2. The EC2 instances are running behind ELB. The admin wants to ensure that the end user request will always go to the EC2 instance where the user session has been created. How can the admin configure this?

- A. Enable ELB cross zone load balancing
- B. Enable ELB cookie setup
- C. Enable ELB sticky session
- D. Enable ELB connection draining

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 98

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 99

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 100

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 101

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 102

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A customer is using AWS for Dev and Test. The customer wants to setup the Dev environment with Cloudformation. Which of the below mentioned steps are not required while using Cloudformation?

- A. Create a stack
- B. Configure a service
- C. Create and upload the template
- D. Provide the parameters configured as part of the template

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 103

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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?

- 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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 104

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 105

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization is generating digital policy files which are required by the admins for verification. Once the files are verified they may not be required in the future unless there is some compliance issue. If the organization wants to save them in a cost effective way, which is the best possible solution?

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 106

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has launched an EBS backed instance. The user started the instance at 9 AM in the morning. Between 9 AM to 10 AM, the user is testing some script. Thus, he stopped the instance twice and restarted it. In the same hour the user rebooted the instance once. For how many instance hours will AWS charge the user?

- A. 3 hours
- B. 4 hours
- C. 2 hours
- D. 1 hour

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 107

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 108

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 109

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 110

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 111

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created a public subnet CIDR (20.0.0.0/24. and VPN only subnets CIDR (20.0.1.0/24. along with the VPN gateway (vgw-12345. to connect to the user's data center. Which of the below mentioned options is a valid entry for the main route table in this scenario?

- A. Destination: 20.0.0.0/24 and Target: vgw-12345
- B. Destination: 20.0.0.0/16 and Target: ALL
- C. Destination: 20.0.1.0/16 and Target: vgw-12345
- D. Destination: 0.0.0.0/0 and Target: vgw-12345

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 112

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 113

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 115

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 116

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 117

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with CIDR 20.0.0.0/16 with only a private subnet and VPN connection using the VPC wizard. The user wants to connect to the instance in a private subnet over SSH. How should the user define the security rule for SSH?

- A. Allow Inbound traffic on port 22 from the user's network
- B. The user has to create an instance in EC2 Classic with an elastic IP and configure the security group of a private subnet to allow SSH from that elastic IP
- C. The user can connect to a instance in a private subnet using the NAT instance
- D. Allow Inbound traffic on port 80 and 22 to allow the user to connect to a private subnet over the Internet

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 118

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created an ELB with the availability zone US-East-1.

The user wants to add more zones to ELB to achieve High Availability. How can the user add more zones to the existing ELB?

- A. It is not possible to add more zones to the existing ELB
- B. The only option is to launch instances in different zones and add to ELB
- C. The user should stop the ELB and add zones and instances as required
- D. The user can add zones on the fly from the AWS console

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 119

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 120

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has configured ELB with two EBS backed EC2 instances. The user is trying to understand the DNS access and IP support for ELB. Which of the below mentioned statements may not help the user understand the IP mechanism supported by ELB?

- A. The client can connect over IPV4 or IPV6 using Dualstack
- B. ELB DNS supports both IPV4 and IPV6
- C. Communication between the load balancer and back-end instances is always through IPV4
- D. The ELB supports either IPV4 or IPV6 but not both

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 121

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 console. The user has to use CLI to provide the specific time
- D. The user can find the data by giving the exact values in the Absolute tab under CloudWatch metrics

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 122

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 123

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

- A. GroupTotalInstances
- B. GroupSumInstances
- C. It is not possible to get a count of all the three metrics together. The user has to find the individual number of running, terminating and pending instances and sum it
- D. GroupInstancesCount

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 124

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 125

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 126

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

- A. The policy allows the IAM user to modify all IAM users' access keys using the console, SDK, CLI or APIs
- B. The policy allows the IAM user to modify all IAM users' credentials using the console, SDK, CLI or APIs
- C. The policy allows the IAM user to modify all credentials using only the console
- D. The policy allows the IAM user to modify the IAM user's own credentials using the console, SDK, CLI or APIs

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 127

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

- A. The access key to connect to the instance is wrong
- B. The security group is not configured properly
- C. The private key used to launch the instance is not correct
- D. The instance CPU is heavily loaded

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 128

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 129

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has configured CloudWatch monitoring on an EBS backed EC2 instance. If the user has not attached any additional device, which of the below mentioned metrics will always show a 0 value?

- A. DiskReadBytes
- B. NetworkIn
- C. NetworkOut
- D. CPUUtilization

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 130

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 131

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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://sqs. 123456789012.us-east-1.amazonaws.com/myqueue`
- D. `http:// 123456789012.sqs. us-east-1.amazonaws.com/myqueue`

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 132

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A sys admin is trying to understand the Auto Scaling activities. Which of the below mentioned processes is not performed by Auto Scaling?

- A. Reboot Instance
- B. Schedule Actions
- C. Replace Unhealthy
- D. Availability Zone Balancing

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 133

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 134

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 135

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 136

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has launched two EBS backed EC2 instances in the US-East-1a region. The user wants to change the zone of one of the instances. How can the user change it?

- A. Stop one of the instances and change the availability zone
- B. The zone can only be modified using the AWS CLI
- C. From the AWS EC2 console, select the Actions - > Change zones and specify new zone
- D. Create an AMI of the running instance and launch the instance in a separate AZ

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 137

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

- A. AWS CloudWatch + AWS SES
- B. AWS CloudWatch + AWS SNS
- C. None. It is not possible to configure the light with the AWS infrastructure services
- D. AWS CloudWatch and a dedicated software turning on the light

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 138

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 139

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 140

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 141

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 142

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 143

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 144

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 145

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 146

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created an ELB with Auto Scaling. Which of the below mentioned offerings from ELB helps the user to stop sending new requests traffic from the load balancer to the EC2 instance when the instance is being deregistered while continuing in-flight requests?

- A. ELB sticky session
- B. ELB deregistration check
- C. ELB connection draining
- D. ELB auto registration Off

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 147

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 148

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 149

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 150

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 151

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 152

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 153

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has launched an EBS backed EC2 instance. The user has rebooted the instance. Which of the below mentioned statements is not true with respect to the reboot action?

- A. The private and public address remains the same
- B. The Elastic IP remains associated with the instance
- C. The volume is preserved
- D. The instance runs on a new host computer

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 154

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has setup a web application on EC2. The user is generating a log of the application performance at every second. There are multiple entries for each second. If the user wants to send that data to CloudWatch every minute, what should he do?

- A. The user should send only the data of the 60th second as CloudWatch will map the receive data timezone with the sent data timezone
- B. It is not possible to send the custom metric to CloudWatch every minute
- C. Give CloudWatch the Min, Max, Sum, and SampleCount of a number of every minute
- D. Calculate the average of one minute and send the data to CloudWatch

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 155

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 156

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

[Show Suggested Answer](#)



Exam question from Amazon's AWS-SysOps

Question #: 157

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 158

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 159

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 160

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a public subnet with VPC and launched an EC2 instance within it. The user is trying to delete the subnet. What will happen in this scenario?

- A. It will delete the subnet and make the EC2 instance as a part of the default subnet
- B. It will not allow the user to delete the subnet until the instances are terminated
- C. It will delete the subnet as well as terminate the instances
- D. The subnet can never be deleted independently, but the user has to delete the VPC first

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 161

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 162

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 credentials is not required while creating the AMI?

- A. AWS account ID
- B. X.509 certificate and private key
- C. AWS login ID to login to the console
- D. Access key and secret access key

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 163

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 164

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 165

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is trying to launch an EBS backed EC2 instance under free usage. The user wants to achieve encryption of the EBS volume. How can the user encrypt the data at rest?

- A. Use AWS EBS encryption to encrypt the data at rest
- B. The user cannot use EBS encryption and has to encrypt the data manually or using a third party tool
- C. The user has to select the encryption enabled flag while launching the EC2 instance
- D. Encryption of volume is not available as a part of the free usage tier

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 166

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 167

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 168

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 169

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has launched an RDS PostgreSQL DB with AWS. The user did not specify the maintenance window during creation. The user has configured RDS to update the DB instance type from micro to large. If the user wants to have it during the maintenance window, what will AWS do?

- A. AWS will not allow to update the DB until the maintenance window is configured
- B. AWS will select the default maintenance window if the user has not provided it
- C. AWS will ask the user to specify the maintenance window during the update
- D. It is not possible to change the DB size from micro to large with RDS

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 170

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console?

- A. The IP address may be attached to one of the instances
- B. The IP address belongs to a different zone than the subnet zone
- C. The user has not created an internet gateway
- D. The IP addresses belong to EC2 Classic; so they cannot be assigned to VPC

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 171

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 172

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with CIDR 20.0.0.0/16 using VPC Wizard. The user has created a public CIDR (20.0.0.0/24) and a VPN only subnet CIDR (20.0.1.0/24) along with the hardware VPN access to connect to the user's data center. Which of the below mentioned components is not present when the VPC is setup with the wizard?

- A. Main route table attached with a VPN only subnet
- B. A NAT instance configured to allow the VPN subnet instances to connect with the internet
- C. Custom route table attached with a public subnet
- D. An internet gateway for a public subnet

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 173

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 174

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 175

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 176

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 177

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is using Cloudformation to launch an EC2 instance and then configure an application after the instance is launched. The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly. How can the user configure this?

- A. It is not possible that the stack creation will wait until one service is created and launched
- B. The user can use the HoldCondition resource to wait for the creation of the other dependent resources
- C. The user can use the DependentCondition resource to hold the creation of the other dependent resources
- D. The user can use the WaitCondition resource to hold the creation of the other dependent resources

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 178

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization has configured two single availability zones. The Auto Scaling groups are configured in separate zones. The user wants to merge the groups such that one group spans across multiple zones. How can the user configure this?

- A. Run the command `as-join-auto-scaling-group` to join the two groups
- B. Run the command `as-update-auto-scaling-group` to configure one group to span across zones and delete the other group
- C. Run the command `as-copy-auto-scaling-group` to join the two groups
- D. Run the command `as-merge-auto-scaling-group` to merge the groups

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 179

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 180

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 181

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has launched two EBS backed EC2 instances in the US-East-1a region. The user wants to change the zone of one of the instances. How can the user change it?

- A. The zone can only be modified using the AWS CLI
- B. Create an AMI of the running instance and launch the instance in a separate AZ
- C. Stop one of the instances and change the availability zone
- D. From the AWS EC2 console, select the Actions - > Change zones and specify the new zone

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 182

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization (account ID 123412341234) has configured the IAM policy to allow the user to modify his credentials. What will the below mentioned statement allow the user to perform?

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "iam:AddUserToGroup",
      "iam:RemoveUserFromGroup",
      "iam:GetGroup"
    ],
    "Resource": "arn:aws:iam:: 123412341234:group/TestingGroup"
  ]
}
```

- A. The IAM policy will throw an error due to an invalid resource name
- B. The IAM policy will allow the user to subscribe to any IAM group
- C. Allow the IAM user to update the membership of the group called TestingGroup
- D. Allow the IAM user to delete the TestingGroup

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 183

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 184

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 185

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered.

What time out period should the user specify for connection draining?

- A. 5 minutes
- B. 1 hour
- C. 30 minutes
- D. 2 hours

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 186

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 187

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 188

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 189

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 190

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 191

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 192

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 193

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 194

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 195

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 196

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 197

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 198

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization has applied the below mentioned policy on an IAM group which has selected the IAM users. What entitlements do the IAM users avail with this policy?

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

- A. The policy is not created correctly. It will throw an error for wrong resource name
- B. The policy is for the group. Thus, the IAM user cannot have any entitlement to this
- C. It allows full access to all AWS services for the IAM users who are a part of this group
- D. If this policy is applied to the EC2 resource, the users of the group will have full access to the EC2 Resources

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 199

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is configuring a CloudWatch alarm on RDS to receive a notification when the CPU utilization of RDS is higher than 50%. The user has setup an alarm when there is some inactivity on RDS, such as RDS unavailability. How can the user configure this?

- A. Setup the notification when the CPU is more than 75% on RDS
- B. Setup the notification when the state is Insufficient Data
- C. Setup the notification when the CPU utilization is less than 10%
- D. It is not possible to setup the alarm on RDS

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 200

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

George has shared an EC2 AMI created in the US East region from his AWS account with Stefano. George copies the same AMI to the US West region. Can Stefano access the copied AMI of George's account from the US West region?

- A. No, copy AMI does not copy the permission
- B. It is not possible to share the AMI with a specific account
- C. Yes, since copy AMI copies all private account sharing permissions
- D. Yes, since copy AMI copies all the permissions attached with the AMI

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 201

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 202

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 203

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 204

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 center
- 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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 205

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user had aggregated the CloudWatch metric data on the AMI ID. The user observed some abnormal behavior 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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 206

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 207

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 SSL protocols is not supported by the security policy?

- A. TLS 1.3
- B. TLS 1.2
- C. SSL 2.0
- D. SSL 3.0

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 208

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with the public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The public subnet uses CIDR 20.0.1.0/24. The user is planning to host a web server in the public subnet (port 80) and a DB server in the private subnet (port 3306). The user is configuring a security group for the public subnet (WebSecGrp) and the private subnet (DBSecGrp). Which of the below mentioned entries is required in the private subnet database security group (DBSecGrp)?

- A. Allow Inbound on port 3306 for Source Web Server Security Group (WebSecGrp)
- B. Allow Inbound on port 3306 from source 20.0.0.0/16
- C. Allow Outbound on port 3306 for Destination Web Server Security Group (WebSecGrp)
- D. Allow Outbound on port 80 for Destination NAT Instance IP

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 209

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created public and VPN only subnets along with hardware VPN access to connect to the user's data center. The user has not yet launched any instance as well as modified or deleted any setup. He wants to delete this VPC from the console. Will the console allow the user to delete the VPC?

- A. Yes, the console will delete all the setups and also delete the virtual private gateway
- B. No, the console will ask the user to manually detach the virtual private gateway first and then allow deleting the VPC
- C. Yes, the console will delete all the setups and detach the virtual private gateway
- D. No, since the NAT instance is running

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 210

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 211

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has setup a custom application which generates a number in decimals. The user wants to track that number and setup the alarm whenever the number is above a certain limit. The application is sending the data to CloudWatch at regular intervals for this purpose. Which of the below mentioned statements is not true with respect to the above scenario?

- A. The user can get the aggregate data of the numbers generated over a minute and send it to CloudWatch
- B. The user has to supply the time zone with each data point
- C. CloudWatch will not truncate the number until it has an exponent larger than 126 (i.e. 1×10^{126})
- D. The user can create a file in the JSON format with the metric name and value and supply it to CloudWatch

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 212

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 behavior is set to Stop
- B. It is not possible to set the termination behavior 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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 213

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 214

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 215

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

1. Perform maintenance on standby
2. Promote standby to primary
3. Perform maintenance on original primary
4. Promote original master back as primary

A. 1, 2, 3, 4

B. 1, 2, 3

C. 2, 3, 1, 4

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 216

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 217

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 218

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created a public subnet CIDR (20.0.0.0/24) and VPN only subnets CIDR (20.0.1.0/24) along with the VPN gateway (vgw-12345) to connect to the user's data center. The user's data center has CIDR 172.28.0.0/12. The user has also setup a NAT instance (i-123456) to allow traffic to the internet from the VPN subnet. Which of the below mentioned options is not a valid entry for the main route table in this scenario?

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 219

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 220

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 221

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 222

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 223

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 224

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 225

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is measuring the CPU utilization of a private data center 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 "data" parameter

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 226

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has enabled detailed CloudWatch monitoring with the AWS Simple Notification Service. Which of the below mentioned statements helps the user understand detailed monitoring better?

- A. SNS will send data every minute after configuration
- B. There is no need to enable since SNS provides data every minute
- C. AWS CloudWatch does not support monitoring for SNS
- D. SNS cannot provide data every minute

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 227

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 228

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 229

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is trying to send custom metrics to CloudWatch using the PutMetricData APIs. Which of the below mentioned points should the user needs to take care while sending the data to CloudWatch?

- A. The size of a request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests
- B. The size of a request is limited to 128KB for HTTP GET requests and 64KB for HTTP POST requests
- C. The size of a request is limited to 40KB for HTTP GET requests and 8KB for HTTP POST requests
- D. The size of a request is limited to 16KB for HTTP GET requests and 80KB for HTTP POST requests

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 230

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An AWS account owner has setup multiple IAM users. One IAM user only has CloudWatch access. He has setup the alarm action which stops the EC2 instances when the CPU utilization is below the threshold limit. What will happen in this case?

- A. It is not possible to stop the instance using the CloudWatch alarm
- B. CloudWatch will stop the instance when the action is executed
- C. The user cannot set an alarm on EC2 since he does not have the permission
- D. The user can setup the action but it will not be executed if the user does not have EC2 rights

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 231

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 232

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a mobile application which makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK and root account access/secret access key to connect to DynamoDB from mobile. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

- A. The user should create a separate IAM user for each mobile application and provide DynamoDB access with it
- B. The user should create an IAM role with DynamoDB and EC2 access. Attach the role with EC2 and route all calls from the mobile through EC2
- C. The application should use an IAM role with web identity federation which validates calls to DynamoDB with identity providers, such as Google, Amazon, and Facebook
- D. Create an IAM Role with DynamoDB access and attach it with the mobile application

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 233

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 HA. Which DB is the user using right now?

- A. My SQL
- B. Oracle
- C. MS SQL
- D. PostgreSQL

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 234

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 235

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 by mistake. The user is trying to create another subnet of CIDR 20.0.0.1/24. How can the user create the second subnet?

- A. There is no need to update the subnet as VPC automatically adjusts the CIDR of the first subnet based on the second subnet's CIDR
- B. The user can modify the first subnet CIDR from the console
- C. It is not possible to create a second subnet as one subnet with the same CIDR as the VPC has been created
- D. The user can modify the first subnet CIDR with AWS CLI

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 236

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with the public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The public subnet uses CIDR 20.0.1.0/24. The user is planning to host a web server in the public subnet (port 80) and a DB server in the private subnet (port 3306). The user is configuring a security group for the public subnet (WebSecGrp) and the private subnet (DBSecGrp). Which of the below mentioned entries is required in the web server security group (WebSecGrp)?

- A. Configure Destination as DB Security group ID (DbSecGrp) for port 3306 Outbound
- B. 80 for Destination 0.0.0.0/0 Outbound
- C. Configure port 3306 for source 20.0.0.0/24 InBound
- D. Configure port 80 InBound for source 20.0.0.0/16

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 237

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 238

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 239

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 240

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is using the AWS SQS to decouple the services. Which of the below mentioned operations is not supported by SQS?

- A. SendMessageBatch
- B. DeleteMessageBatch
- C. CreateQueue
- D. DeleteMessageQueue

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 241

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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`

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 242

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 243

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 244

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with CIDR 20.0.0.0/24. The user has used all the IPs of CIDR and wants to increase the size of the VPC. The user has two subnets: public (20.0.0.0/28) and private (20.0.1.0/28). How can the user change the size of the VPC?

- A. The user can delete all the instances of the subnet. Change the size of the subnets to 20.0.0.0/32 and 20.0.1.0/32, respectively. Then the user can increase the size of the VPC using CLI
- B. It is not possible to change the size of the VPC once it has been created
- C. The user can add a subnet with a higher range so that it will automatically increase the size of the VPC
- D. The user can delete the subnets first and then modify the size of the VPC

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 245

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 246

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 247

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has configured an ELB to distribute the traffic among multiple instances. The user instances are facing some issues due to the back-end servers. Which of the below mentioned CloudWatch metrics helps the user understand the issue with the instances?

- A. HTTPCode_Backend_3XX
- B. HTTPCode_Backend_4XX
- C. HTTPCode_Backend_2XX
- D. HTTPCode_Backend_5XX

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 248

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 249

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user runs the command `dd if=/dev/zero of=/dev/xvdfbs=1M` on a fresh blank EBS volume attached to a Linux instance. Which of the below mentioned activities is the user performing with the command given above?

- A. Creating a file system on the EBS volume
- B. Mounting the device to the instance
- C. Pre warming the EBS volume
- D. Formatting the EBS volume

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 250

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 251

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 252

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 253

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with public and private subnets. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.1.0/24 and the public subnet uses CIDR 20.0.0.0/24. The user is planning to host a web server in the public subnet (port 80) and a DB server in the private subnet (port 3306). The user is configuring a security group of the NAT instance. Which of the below mentioned entries is not required for the NAT security group?

- A. For Inbound allow Source: 20.0.1.0/24 on port 80
- B. For Outbound allow Destination: 0.0.0.0/0 on port 80
- C. For Inbound allow Source: 20.0.0.0/24 on port 80
- D. For Outbound allow Destination: 0.0.0.0/0 on port 443

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 254

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 255

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowUsersAllActionsForCredentials",
    "Effect": "Allow",
    "Action": [
      "iam:*LoginProfile",
      "iam:*AccessKey*",
      "iam:*SigningCertificate*"
    ],
    "Resource": ["arn:aws:iam:: 123412341234:user/${aws:username}"]
  }]
}
```

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 256

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A sys admin is trying to understand the sticky session algorithm. Please select the correct sequence of steps, both when the cookie is present and when it is not, to help the admin understand the implementation of the sticky session:

1. ELB inserts the cookie in the response
2. ELB chooses the instance based on the load balancing algorithm
3. Check the cookie in the service request
4. The cookie is found in the request
5. The cookie is not found in the request

A. 3,1,4,2 [Cookie is not Present] & 3,1,5,2 [Cookie is Present]

B. 3,4,1,2 [Cookie is not Present] & 3,5,1,2 [Cookie is Present]

C. 3,5,2,1 [Cookie is not Present] & 3,4,2,1 [Cookie is Present]

D. 3,2,5,4 [Cookie is not Present] & 3,2,4,5 [Cookie is Present]

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 257

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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. Time zone

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 258

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 259

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is trying to connect to a running EC2 instance using SSH. However, the user gets an Unprotected Private Key File error. Which of the below mentioned options can be a possible reason for rejection?

- A. The private key file has the wrong file permission
- B. The ppk file used for SSH is read only
- C. The public key file has the wrong permission
- D. The user has provided the wrong user name for the OS login

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 260

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 261

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 262

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 263

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has scheduled the maintenance window of an RDS DB on Monday at 3 AM. Which of the below mentioned events may force to take the DB instance offline during the maintenance window?

- A. Enabling Read Replica
- B. Making the DB Multi AZ
- C. DB password change
- D. Security patching

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 264

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 265

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has configured Auto Scaling with the minimum capacity as 2 and the desired capacity as 2. The user is trying to terminate one of the existing instance with the command:

```
as-terminate-instance-in-auto-scaling-group<Instance ID>--decrement-  
desired-capacity
```

What will Auto Scaling do in this scenario?

- A. Terminates the instance and does not launch a new instance
- B. Terminates the instance and updates the desired capacity to 1
- C. Terminates the instance and updates the desired capacity and minimum size to 1
- D. Throws an error

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 266

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 267

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 268

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 269

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is having data generated randomly based on a certain event. The user wants to upload that data to CloudWatch. It may happen that event may not have data generated for some period due to randomness. Which of the below mentioned options is a recommended option for this case?

- A. For the period when there is no data, the user should not send the data at all
- B. For the period when there is no data the user should send a blank value
- C. For the period when there is no data the user should send the value as 0
- D. The user must upload the data to CloudWatch as having no data for some period will cause an error at CloudWatch monitoring

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 270

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 271

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 272

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 273

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 274

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 275

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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": "*"

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 276

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 277

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 278

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 279

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has two EC2 instances running in two separate regions. The user is running an internal memory management tool, which captures the data and sends it to CloudWatch in US East, using a CLI with the same namespace and metric. Which of the below mentioned options is true with respect to the above statement?

- A. The setup will not work as CloudWatch cannot receive data across regions
- B. CloudWatch will receive and aggregate the data based on the namespace and metric
- C. CloudWatch will give an error since the data will conflict due to two sources
- D. CloudWatch will take the data of the server, which sends the data first

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 280

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization has created a Queue named `modularqueue` with SQS. The organization is not performing any operations such as SendMessage, ReceiveMessage, DeleteMessage, GetQueueAttributes, SetQueueAttributes, AddPermission, and RemovePermission on the queue. What can happen in this scenario?

- A. AWS SQS sends notification after 15 days for inactivity on queue
- B. AWS SQS can delete queue after 30 days without notification
- C. AWS SQS marks queue inactive after 30 days
- D. AWS SQS notifies the user after 2 weeks and deletes the queue after 3 weeks.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 281

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization has setup Auto Scaling with ELB. Due to some manual error, one of the instances got rebooted. Thus, it failed the Auto Scaling health check. Auto Scaling has marked it for replacement. How can the system admin ensure that the instance does not get terminated?

- A. Update the Auto Scaling group to ignore the instance reboot event
- B. It is not possible to change the status once it is marked for replacement
- C. Manually add that instance to the Auto Scaling group after reboot to avoid replacement
- D. Change the health of the instance to healthy using the Auto Scaling commands

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 282

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 283

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 284

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 285

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a VPC with two subnets: one public and one private. The user is planning to run the patch update for the instances in the private subnet. How can the instances in the private subnet connect to the Internet?

- A. Use the internet gateway with a private IP
- B. Allow outbound traffic in the security group for port 80 to allow internet updates
- C. The private subnet can never connect to the internet
- D. Use NAT with an elastic IP

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 286

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has configured an EC2 instance in the US-East-1a zone. The user has enabled detailed monitoring of the instance. The user is trying to get the data from CloudWatch using a CLI. Which of the below mentioned CloudWatch endpoint URLs should the user use?

- A. monitoring.us-east-1.amazonaws.com
- B. monitoring.us-east-1-a.amazonaws.com
- C. monitoring.us-east-1a.amazonaws.com
- D. cloudwatch.us-east-1a.amazonaws.com

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 287

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has configured ELB with Auto Scaling. The user suspended the Auto Scaling AddToLoadBalancer (which adds instances to the load balancer. process for a while). What will happen to the instances launched during the suspension period?

- A. The instances will not be registered with ELB and the user has to manually register when the process is resumed
- B. The instances will be registered with ELB only once the process has resumed
- C. Auto Scaling will not launch the instance during this period due to process suspension
- D. It is not possible to suspend only the AddToLoadBalancer process

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 288

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 289

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has moved an object to Glacier using the life cycle rules. The user requests to restore the archive after 6 months. When the restore request is completed the user accesses that archive. Which of the below mentioned statements is not true in this condition?

- A. The archive will be available as an object for the duration specified by the user during the restoration request
- B. The restored object's storage class will be RRS
- C. The user can modify the restoration period only by issuing a new restore request with the updated period
- D. The user needs to pay storage for both RRS (restored) and Glacier (Archive. Rates)

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 290

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is running a batch process on EBS backed EC2 instances. The batch process starts a few instances to process hadoop Map reduce jobs which can run between 50 ~ 600 minutes or sometimes for more time. The user wants to configure that the instance gets terminated only when the process is completed. How can the user configure this with CloudWatch?

- A. Setup the CloudWatch action to terminate the instance when the CPU utilization is less than 5%
- B. Setup the CloudWatch with Auto Scaling to terminate all the instances
- C. Setup a job which terminates all instances after 600 minutes
- D. It is not possible to terminate instances automatically

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 291

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 292

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 293

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 294

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 295

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

[1]

policy define?

```
"Statement": [{  
  "Sid": "Stmt1388811069831",  
  "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 [1] [1]

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 296

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 297

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has created a launch configuration for Auto Scaling where CloudWatch detailed monitoring is disabled. The user wants to now enable detailed monitoring. How can the user achieve this?

- A. Update the Launch config with CLI to set InstanceMonitoringDisabled = false
- B. The user should change the Auto Scaling group from the AWS console to enable detailed monitoring
- C. Update the Launch config with CLI to set InstanceMonitoring.Enabled = true
- D. Create a new Launch Config with detail monitoring enabled and update the Auto Scaling group

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 298

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 299

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 300

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 301

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has enabled termination protection on an EC2 instance. The user has also set Instance initiated shutdown behavior to terminate. When the user shuts down the instance from the OS, what will happen?

- A. The OS will shutdown but the instance will not be terminated due to protection
- B. It will terminate the instance
- C. It will not allow the user to shutdown the instance from the OS
- D. It is not possible to set the termination protection when an Instance initiated shutdown is set to Terminate

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 302

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 303

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has launched a Windows based EC2 instance. However, the instance has some issues and the user wants to check the log. When the user checks the Instance console output from the AWS console, what will it display?

- A. All the event logs since instance boot
- B. The last 10 system event log error
- C. The Windows instance does not support the console output
- D. The last three system events' log errors

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 304

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 192.166 100.0 full access to the "mybucket" bucket
- B. Denies the server with the IP address 192.166 100.188 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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 305

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which services allow the customer to retain run administrative privileges on the underlying EC2 instances? (Choose two.)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 306

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 307

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 308

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

- A. AWS Direct Connect
- B. Placement Groups
- C. VPC private subnets
- D. EC2 Dedicated Instances
- E. Multiple Availability Zones

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 309

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Amazon EBS snapshots have which of the following two characteristics? (Choose two.)

- A. EBS snapshots only save incremental changes from snapshot to snapshot
- B. EBS snapshots can be created in real-time without stopping an EC2 instance
- C. EBS snapshots can only be restored to an EBS volume of the same size or smaller
- D. EBS snapshots can only be restored and mounted to an instance in the same Availability Zone as the original EBS volume

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 310

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 311

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You run a web application with the following components Elastic Load Balancer (ELB), 3 Web/Application servers, 1 MySQL RDS database with read replicas, and Amazon Simple Storage Service (Amazon S3) for static content. Average response time for users is increasing slowly.

What three CloudWatch RDS metrics will allow you to identify if the database is the bottleneck? (Choose three.)

- A. The number of outstanding IOs waiting to access the disk.
- B. The amount of write latency.
- C. The amount of disk space occupied by binary logs on the master.
- D. The amount of time a Read Replica DB Instance lags behind the source DB Instance
- E. The average number of disk I/O operations per second.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 312

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 313

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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 two.)

- 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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 314

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 315

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 316

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

When you put objects in Amazon S3, what is the indication that an object was successfully stored?

- A. Each S3 account has a special bucket named_s3_logs. Success codes are written to this bucket with a timestamp and checksum.
- B. A success code is inserted into the S3 object metadata.
- C. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful.
- D. Amazon S3 is engineered for 99.999999999% durability. Therefore, there is no need to confirm that data was inserted.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 317

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 318

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You have a business-to-business web application running in a VPC consisting of an Elastic Load Balancer (ELB), web servers, application servers and a database. Your web application should only accept traffic from pre-defined customer IP addresses.

Which two options meet this security requirement? (Choose two.)

- A. Configure web server VPC security groups to allow traffic from your customers' IPs
- B. Configure your web servers to filter traffic based on the ELB's "X-forwarded-for" header
- C. Configure ELB security groups to allow traffic from your customers' IPs and deny all outbound traffic
- D. Configure a VPC NACL to allow web traffic from your customers' IPs and deny all outbound traffic

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 319

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 320

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 321

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 322

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In AWS, which security aspects are the customer's responsibility? (Choose four.)

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 323

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 324

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 325

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Your mission is to create a lights-out datacenter environment, and you plan to use AWS OpsWorks to accomplish this. First you created a stack and added an App Server layer with an instance running in it. Next you added an application to the instance, and now you need to deploy a MySQL RDS database instance. Which of the following answers accurately describe how to add a backend database server to an OpsWorks stack? (Choose three.)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 326

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user needs to put sensitive data in an Amazon S3 bucket that can be accessed through an S3 VPC endpoint only. The user must ensure that resources in the VPC can only access the single S3 bucket.

Which combination of actions will meet the requirements? (Choose two.)

- A. Configure the bucket policy to only allow access through the S3 Private Endpoint.
- B. Modify the VPC endpoint policy on the bucket to only allow the VPC to access it.
- C. Modify the VPC peering configuration to only allow access to the S3 private Endpoint.
- D. Configure the VPC endpoint policy to only allow the VPC to access the specific S3 bucket.
- E. Configure the IAM policy attached to the S3 bucket to only allow access from the specific VPC.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 327

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A corporate website is hosted on several Amazon EC2 instances across multiple regions around the globe.

How should an Administrator configure the website to maintain high availability with minimal downtime if one of the regions has network connectivity congestion for an extended period of time?

- A. Create an Elastic Load Balancer in front of all the Amazon EC2 instances.
- B. Create an Elastic Load Balancer that fails over to the secondary site when the primary site is not reachable.
- C. Create an Amazon Route 53 Latency Based Routing Record Set that resolves to an Elastic Load Balancer in each region. Set an appropriate health check on each ELB.
- D. Create an Amazon Route 53 latency Based Routing Record Set that resolves to Elastic Load Balancers in each region and has the Evaluate Target Health flag set to true.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 328

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A database running on Amazon EC2 requires sustained IOPS performance.

Which kind of Amazon EBS volume should an Administrator choose for this solution?

- A. Cloud HDD
- B. General Purpose SSD
- C. Provisioned IOPS SSD
- D. Throughput Optimized HDD

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 329

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does the `configure` command allow an Administrator to do when setting up the AWS CLI? (Choose two.)

- A. Decide which VPC to create instances in.
- B. Designate the format of the response to CLI commands.
- C. Choose the default EC2 instance.
- D. Encrypt the CLI commands.
- E. Designate the default region.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 330

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An Administrator has an Amazon EC2 instance with an IPv6 address. The Administrator needs to prevent direct access to this instance from the Internet. The Administrator should place the EC2 instance in a:

- A. Private Subnet with an egress-only Internet Gateway attached to the subnet and placed in the subnet Route Table.
- B. Public subnet with an egress-only Internet Gateway attached to the VPC and placed in the VPC Route Table.
- C. Private subnet with an egress-only Internet Gateway attached to the VPC and placed in the subnet Route Table.
- D. Public subnet and a security group that blocks inbound IPv6 traffic attached to the interface.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 331

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

As part of an operational audit, an Administrator is tasked with showing that all security responsibilities under the customer's control are properly executed. Which of the following items is the customer responsible for providing to the auditor? (Choose two.)

- A. Physical data center access logs
- B. AWS CloudTrail logs showing API calls
- C. Amazon EC2 instance system logs
- D. Storage device destruction records
- E. Xen Hypervisor system logs

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 332

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A colleague is attempting to launch several new CloudFormation stacks, and receives the following error response:

```
<Error>  
  <Type>Sender</Type>  
  <Code>Throttling</Code>  
  <Message>Rate exceeded</Message>  
</Error>
```

What should be done to address the error?

- A. Add a Pause to the CloudFormation templates.
- B. Add an exponential backoff between CreateStack API calls.
- C. Run the CloudFormation API calls from a larger Amazon EC2 instance.
- D. Combine stack templates into one, and retry the CreateStack API call.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 333

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A security policy allows instances in the Production and Development accounts to write application logs to an Amazon S3 bucket belonging to the Security team's account. Only the Security team should be allowed to delete logs from the S3 bucket.

Using the `myAppRole` EC2 role, the production and development teams report that the application servers are not able to write to the S3 bucket.

Which changes need to be made to the policy to allow the application logs to be written to the S3 bucket?

Production Account: 111111111111

Dev Account: 222222222222 -

Security Account: 555555555555 -

```
{
  "Version": "2012-10-17",
  "Statement": [ [
    "Effect": "Allow",
    "Principal": [ {
      "AWS": [
        "arn: aws:iam: : 111111111111: role/myAppRole"
        "arn: aws:iam: : 222222222222: role/myAppRole"
      ]
    } ],
    "Action": [
      "s3: *"
    ],
    "Resource": [
    ],
    "Condition" {
      "StringNotLike": {
        "aws: userID": [
          "555555555555"
        ]
      }
    }
  ] ]
}
```

- A. Update the Action for the Allow policy from `s3:*` to `s3:PutObject`
- B. Change the order of the statements in the bucket policy, moving the Deny policy above the Allow policy.
- C. Update the Action for the Deny policy from `s3:*` to `s3: Delete*`.
- D. Remove the bucket policy, because the default security behavior will not allow objects to be deleted by non bucket owners.

Show Suggested Answer

Exam question from Amazon's AWS-SysOps

Question #: 334

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company is auditing their infrastructure to obtain a compliance certification.

Which of the following options are the company's responsibility within the Shared Responsibility Model? (Choose two.)

- A. AWS API endpoint SSL Certificates
- B. EC2 Instance Operating System updates
- C. EBS Encryption-at-rest algorithms
- D. IAM user password policies
- E. AWS Hypervisor software updates

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 335

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which instance characteristics are required if an Administrator wants to ensure use of the Amazon EC2 auto-recovery option? (Choose two.)

- A. The instance only has EBS volumes.
- B. The instance has EC2 Instance Store root volumes.
- C. The tenancy attribute is set to `default` (shred tenancy).
- D. The tenancy attribute is set to `Dedicated`.
- E. The instance type belongs to the d2, i2 or i3 instance type.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 336

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which two steps are required to generate a report detailing specific cost allocation tags when creating a Monthly Cost Allocation report? (Choose two.)

- A. Use AWS CloudTrail to export the events for the specified resources.
- B. Use an AWS Lambda function to read the resources' metadata, and write the specified tags to a DynamoDB table.
- C. Activate the `requested` tags by clicking Manage report tags on the Billing Preferences page.
- D. Select the checkbox for Cost Allocation Report in the AWS account's Billing Management Console.
- E. Create a new Budget using the Billing Management Console, use the `Include costs related to Tags` feature, and select the requested tags.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 337

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company has a fleet of EC2 instances, and needs to remotely execute scripts for all of the instances.

Which Amazon EC2 Systems Manager feature allows this?

- A. System Manager Automation
- B. System Manager Run Command
- C. System Manager Parameter Store
- D. System Manager Inventory

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 338

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A corporate policy requires all new infrastructure deployments to use scalable and reusable resources to improve resources delivery times. The policy also restricts resource configuration management to the systems operations team. The development team requests the ability to deploy resources on demand in an effort to streamline their software development lifecycle.

What can the systems operations team do to ensure company policy is followed while also meeting the development team's requests?

- A. Create an AWS CloudFormation on template with the requested resources, and give it to the development team to adjust as needed.
- B. Provision the resources using the CLI, and create the necessary IAM permissions to allow the development team to modify them as needed.
- C. Create the AWS Service Catalog product and share with the development team through the Service Catalog.
- D. Grant the development team access to the AWS CloudFormation Design Template Editor to specify the needed resources and configurations. Once the templates are complete, the system operations team will launch the resources.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 339

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An application hosted on AWS is going through an external compliance assessment. An Administrator has been tasked with providing proof of physical security at the facilities that are hosting the application.

What should the Administrator do?

- A. Work with AWS support to schedule a tour for the auditors.
- B. Send a copy of the AWS Security whitepaper to the auditors.
- C. Obtain a relevant report from AWS Artifact and share it with the auditors.
- D. Find the address for the AWS Direct Connect facility on the AWS Website.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 340

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What can an Administrator do to monitor whether an organization's instances are compliant with corporate policies and guidelines?

- A. Check the instances' metadata to determine what software is running.
- B. Use AWS CloudTrail logs to identify the applications running on the instances.
- C. Set CloudWatch alarms that are triggered with any software change on the instances.
- D. Using Config Rules in the AWS Config service to check the instance's configuration and applications.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 341

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following are the customer's responsibilities, according to the AWS Shared Responsibility Security Model? (Choose two.)

- A. Operating system, network, and firewall configuration
- B. Client-side data encryption and data integrity authentication
- C. AWS data center access logs
- D. Hypervisor updates and configuration
- E. Physical media destruction

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 342

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Elastic Load Balancing automatically distributes incoming traffic across multiple ____ instances.

- A. EC2
- B. RDS
- C. M3
- D. DB

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 343

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

_____ is a task coordination and state management service for cloud applications.

- A. Amazon SWF
- B. Amazon FPS
- C. Amazon SES
- D. Amazon SNS

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 344

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A block device is a storage device that moves data in sequences. How many types of block devices does Amazon EC2 support?

- A. 2 -instance store volumes and EBS volumes
- B. 5 -General Purpose SSD, Provisioned IOPS SSD, Throughput Optimized HDD, Cold HDD, and Magnetic
- C. 3 -SSD, HDD, and Magnetic
- D. 1 -instance store volumes

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 345

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Do Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

- A. No, they are dependent.
- B. No, you cannot attach EBS volumes to an instance.
- C. Yes, they do but only if they are detached from the instance.
- D. Yes, they do, if the Delete on termination flag is unset.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 346

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Is it possible to access S3 objects from the Internet?

- A. Yes, but it has to pass through EC2.
- B. Yes, it is possible if proper public readable accesses and ACLs are set.
- C. No, there is no way to access any S3 objects from the Internet.
- D. No, only a general overview of S3 objects can be read from the Internet.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 347

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

_____ is a fast, reliable, scalable, fully managed message queuing service.

- A. AWS Data Pipeline
- B. Amazon SES
- C. Amazon SQS
- D. Amazon SNS

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 348

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon Route53 provide?

- A. A global Content Delivery Network
- B. A scalable DNS web service
- C. An SSH endpoint for Amazon EC2
- D. None of these

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 349

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon VPC stand for?

- A. Amazon Virtual Private Cloud
- B. Amazon Variable Power Cluster
- C. Amazon Virtual Private Computer
- D. Amazon Virtual Public Cloud

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 350

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following does Amazon S3 provide?

- A. A virtual server in the cloud
- B. A highly-scalable cloud storage
- C. A highly encrypted virtual disk in the cloud
- D. A transient storage in the cloud

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 351

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The billing process for Amazon EC2 instances was updated as of October 2, 2017. Which of the following statements is true regarding how you pay for Amazon EC2 instances? (Choose two.)

- A. Payment does not vary based on the instance AMI's operating system.
- B. You can pay per hour or per second, depending on the instance AMI's operating system.
- C. You pay for compute capacity by the day; hours are billed in proportion.
- D. You can pay per hour or per second, depending on the instance type.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 352

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

When an instance terminates, Amazon EC2 uses the value of the _____ attribute for each attached Amazon EBS volume to determine whether to preserve or delete the volume.

- A. InstanceInitiatedShutdownBehavior
- B. DeleteOnTermination
- C. EC2ModifyInstance
- D. DisableApiTermination

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 353

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon RDS perform?

- A. It tests the functionalities in websites.
- B. It blocks users from creating DB instances.
- C. It manages the work involved in setting up a relational database.
- D. It provides sensory feedback.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 354

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What was the recommended use case for S3 Reduced Redundancy storage before its deprecation was planned?

- A. It was used to reduce storage costs by providing 500 times the durability of a typical disk drive at lower levels of redundancy.
- B. It was used to reduce storage costs for noncritical data at lower levels of redundancy.
- C. It was used to reduce storage costs by allowing you to destroy any copy of your files outside a specific jurisdiction.
- D. It was used to reduce storage costs for reproducible data at high levels of redundancy in a single facility.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 355

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

----- is a fast, flexible, fully managed pub/sub messaging service.

- A. Amazon SQS
- B. Amazon SES
- C. Amazon FPS
- D. Amazon SNS

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 356

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Does AWS offer any web-based graphic user interface to access and manage EC2 instances?

- A. Yes, the AWS Application Clusters.
- B. No, you can only use the available software development kits.
- C. Yes, the AWS Management Console.
- D. No, you can only use the command line interface.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 357

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is the maximum size of an object in Amazon S3?

- A. 4 TB
- B. Unlimited
- C. 5 TB
- D. 500 MB

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 358

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Amazon EBS provides the ability to create backups of any Amazon EC2 volume into what is known as ____.

- A. snapshots
- B. mirrors
- C. instance backups
- D. images

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 359

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following size ranges is true of Individual Amazon S3 objects?

- A. 5 gigabytes to 5 terabytes
- B. 0 bytes to 5 terabytes
- C. 100 megabytes to 5 gigabytes
- D. 1 byte to 5 gigabytes

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 360

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is a security group in Amazon AWS?

- A. A UNIX Group that gives permission to edit security settings
- B. An authorized group of instances that control access to other resources
- C. A virtual firewall that controls the traffic for one or more instances
- D. An Access Control List (ACL) for AWS resources

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 361

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon EBS stand for?

- A. Elastic Business Server
- B. Elastic Basic Storage
- C. Elastic Blade Server
- D. Elastic Block Store

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 362

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

EBS (Elastic Block Store) can be best described as:

- A. persistent internet storage.
- B. persistent block storage.
- C. transient instance storage.
- D. transient block storage.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 363

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In Amazon RDS, which of the following provides enhanced availability and durability for Database (DB) Instances, making them to be a natural fit for production database workloads?

- A. Placement Groups
- B. Multi-Option Group deployment
- C. Multi-AZ deployment
- D. Multi-VPC deployment

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 364

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The Amazon Linux AMI is:

- A. a simple OS installation media.
- B. an instance package provided by the AWS.
- C. a refined, easy-to-use, up-to-date Linux desktop distribution.
- D. a supported and maintained Linux image provided by AWS.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 365

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

_____ is a fully managed service for real-time processing of streaming data at massive scale.

- A. AWS Data Pipeline
- B. Amazon Kinesis
- C. AWS CloudHSM
- D. Amazon Elastic Compute Cloud

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 366

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In Amazon S3, what is the document that defines who can access a particular bucket or object called?

- A. Access Control Record
- B. Access Control Service
- C. Access Control List
- D. Access Control Server

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 367

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Where is an object stored in Amazon S3?

- A. in a Bucket
- B. in a Collector
- C. in an Archive
- D. in a Vault

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 368

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which AWS service offers cost optimization by launching instances automatically only when need-ed?

- A. Elastic Load Balancing
- B. Elastic Compute Cloud
- C. Auto Scaling
- D. Relational Database Service

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 369

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon SWF stand for?

- A. Simple Waveflow Service
- B. Simple WebFactor Service
- C. Simple Workflow Service
- D. Simple WebForm Service

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 370

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In Amazon EC2, can you create an EBS volume from a snapshot and attach it to another instance?

- A. No, you cannot attach EBS volumes to an instance.
- B. Yes, you can but only if the volume is larger than 2TB.
- C. No, you can't create an EBS volume from a snapshot.
- D. Yes, you can.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 371

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Spot instances are ideally designed for which purpose below?

- A. Running database instances that can scale up and down based on a specific workload.
- B. Running long duration and highly transactional applications.
- C. For building distributed fault tolerant databases under a tight deadline.
- D. Taking advantage of excess EC2 capacity at prices below standard on-demand rates, for short duration jobs.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 372

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon EMR stand for?

- A. Elastic Magnetic Resonance
- B. Encrypted Machine Reads
- C. Elastic MapReduce
- D. Encrypted Machine Rendering

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 373

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is the main use of EMR?

- A. Data-sensitive storage
- B. Encryption
- C. Data-intensive processing tasks
- D. authentication

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 374

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What cloud service does Amazon S3 offer?

- A. Atomic updates across keys over the Internet
- B. Messaging over the Internet
- C. Storage over the Internet
- D. Object locking over the Internet

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 375

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is launching an instance with EC2. Which options below should the user consider before launching an instance?

- A. Select the region where the instance is being launched.
- B. All choices are correct.
- C. Select the instance type.
- D. Select the OS of the AMI.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 376

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In regard to AWS CloudFormation, to pass values to your template at runtime you should use _____.

- A. parameters
- B. conditions
- C. resources
- D. mapping

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 377

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon S3 stand for?

- A. Social Storage Service
- B. Simple Storage Service
- C. Secure Storage Service
- D. Standard Storage Service

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 378

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is Amazon WorkSpaces?

- A. Amazon WorkSpaces is a fully managed desktop computing service in the cloud, allowing end-users to access the documents, applications, and resources they need with the device of their choice.
- B. Amazon WorkSpaces is a flexible application management solution with automation tools that enable you to model and control your applications and their supporting infrastructure.
- C. Amazon WorkSpaces is a fully redundant data storage infrastructure for storing and retrieving any amount of data, at any time, from anywhere on the web.
- D. Amazon WorkSpaces is a web service that enables businesses, researchers, data analysts, and de-velopers to easily and cost-effectively process vast amounts of data.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 379

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does AMI stand for?

- A. Amazon Machine Image
- B. Advanced Machine Instance
- C. Amazon Micro Instance
- D. Advanced Machine Image

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 380

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following statements is true of tags and resource identifiers for EC2 instances?

- A. You can't select instances by their tags for stoppage, termination, or deletion
- B. You don't need to specify the resource identifier while terminating a resource.
- C. You don't need to specify the resource identifier while stopping a resource.
- D. You can select instances by their tags for stoppage, termination, or deletion

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 381

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon RDS stand for?

- A. Amazon Regional Data Server
- B. Amazon Regional Database Service
- C. Amazon Relative Data Service
- D. Amazon Relational Database Service

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 382

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon SES provide?

- A. A managed Email Server
- B. A scalable anti-spam service
- C. A scalable email sending and receiving service
- D. A managed drag-and-drop interface with the AWS CloudFormation Designer

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 383

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Pricing is _____ consumed for EC2 instances.

- A. per instance-hour only
- B. per instance-minute or instance-hour
- C. per instance-second or per instance-hour
- D. per instance-minute only

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 384

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon SES stand for?

- A. Simple Elastic Server
- B. Software Email Solution
- C. Software Enabled Server
- D. Simple Email Service

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 385

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is a "vault" in Amazon Glacier?

- A. A unique ID that maps an AWS Region, plus a specific Amazon S3 bucket
- B. A way to group archives together in Amazon Glacier
- C. A container for storing S3 buckets
- D. A free tier available for 12 months following your AWS sign-up date

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 386

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has launched five instances and have registered them with an ELB. How can the user add the sixth EC2 instance to the ELB?

- A. The user must stop the ELB and add the sixth instance.
- B. The user can add the sixth instance on the fly through API, CLI or the AWS Management Con-sole.
- C. The user can add the instance and change the ELB config file.
- D. The ELB can only have a maximum of five instances.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 387

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following programming languages is not supported by Amazon's Elastic Beanstalk?

- A. Ruby
- B. Java
- C. Node.js
- D. Perl

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 388

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Amazon CloudFront is a _____.

- A. persistent block level storage volume
- B. content delivery network service
- C. fully managed desktop computing service in the cloud
- D. task coordination and state management service for cloud applications

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 389

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In EC2, what happens to the data in an instance store if an instance reboots (either intentionally or unintentionally)?

- A. Data is partially present in the instance store.
- B. Data persists in the instance store.
- C. Data is deleted from the instance store for security reasons.
- D. Data in the instance store will be lost.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 390

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon EC2 provide?

- A. A platform to run code (Java, PHP, Python), paying on an hourly basis
- B. A physical computing environment
- C. Virtual Server Hosting
- D. Domain Name System (DNS)

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 391

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does RRS stand, in the context of S3 services?

- A. Regional Rights Storage
- B. Relational Rights Storage
- C. Regional Rights Standard
- D. Reduced Redundancy Storage

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 392

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which choice is a storage option supported by Amazon EC2?

- A. Amazon SNS store
- B. Amazon Instance Store
- C. Amazon AppStream store
- D. None of these

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 393

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Amazon EC2 provides virtual computing environments known as _____.

- A. instances
- B. volumes
- C. microsystems
- D. servers

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 394

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following services is offered by CloudWatch?

- A. Fixing broken links on the client's instances
- B. Creating IAM users for all services in AWS
- C. Monitoring estimated AWS charges
- D. Balancing the request load between various instances

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 395

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The fastest way to load 300 TB of data to AWS is _____.

- A. to directly upload all data to S3 over a dedicated 100 Mbps connection
- B. to use AWS Import/Export Snowball
- C. to use VM Import/Export
- D. to zip all the data and then upload to S3

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 396

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

AMIs can be _____.

- A. only private unless created by Amazon
- B. created only by Amazon
- C. created only for Linux instances
- D. public or private

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 397

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Can you configure multiple Load Balancers with a single Auto Scaling group?

- A. Yes, you can provide the ELB is configured with Amazon AppStream.
- B. No
- C. Yes
- D. Yes, you can but only if it is configured with Amazon Redshift.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 398

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is sending custom data metrics to CloudWatch. What is the allowed time stamp granularity for each data point published for the custom metric?

- A. 1 nanosecond
- B. 1 millisecond
- C. 1 minute
- D. 1 second

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 399

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

When rebalancing, Auto Scaling launches new instances before terminating the old ones, so that re-balancing does not compromise the performance or availability of your application. Because Auto Scaling attempts to launch new instances before terminating the old ones, being at or near the specified maximum capacity could impede or completely halt rebalancing activities. What does Auto Scaling do in order to avoid this problem?

- A. It can temporarily exceed the specified maximum capacity of a group by a 20 percent margin (or by a 2-instance margin, whichever is greater) during a rebalancing activity.
- B. It can add new reserved instances you have defined.
- C. It can temporarily exceed the specified maximum capacity of a group by a 10 percent margin (or by a 1-instance margin, whichever is greater) during a rebalancing activity.
- D. It can temporarily exceed the specified maximum capacity of a group by a 5 percent margin (or by a 1-instance margin, whichever is greater) during a rebalancing activity.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 400

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does the AWS Storage Gateway provide?

- A. It provides data security features by enabling an encrypted data storage on Amazon S3.
- B. It provides an encrypted SSL endpoint for backups in the cloud.
- C. It provides seamless integration with data security features between your on-premises IT environment and the Amazon Web Services (AWS) storage infrastructure.
- D. It provides a backup solution to on-premises Cloud storage.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 401

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

By default, how many Elastic IP addresses can you have per region for your EC2 instances?

- A. 10
- B. 2
- C. 20
- D. 5

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 402

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Elasticity is one of the benefits of using Elastic Beanstalk. Which of the following best describes the concept of elasticity?

- A. It is the ability for counting the number of architectural design considerations that are required to develop a console.
- B. It is the streamlining of resource acquisition and release, so that your infrastructure can rapidly scale in and scale out as demand fluctuates.
- C. It is the process of examining the amount of security credentials required to access a data volume.
- D. It is the procedure of estimating the resource cost, so that you can run a specific project on AWS.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 403

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is an Auto Scaling group?

- A. It is a group of ELBs that are used to add instances from various regions.
- B. It is a logical grouping of EC2 instances that share similar characteristics for scaling and management.
- C. It is a collection of EC2 instance launch parameters with different characteristics for scaling and management.
- D. It is a group of launch configurations for Elastic load balancers in the same region.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 404

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which service is offered by Auto Scaling?

- A. Automatic scaling storage
- B. Automatic scale EC2 capacity
- C. Automatic scale ECS capacity
- D. Automatic scale elastic IP

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 405

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the scaling options given below is not supported by Auto Scaling?

- A. All these options are supported by Auto Scaling
- B. Manual scaling
- C. Scaling based on CPU utilization
- D. Scaling based on time

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 406

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Security groups in Amazon VPC _____.

- A. control incoming traffic only
- B. control both inbound and outbound traffic
- C. control neither incoming nor outgoing traffic
- D. control outgoing traffic only

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 407

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

_____ in VPC are stateful where return traffic is automatically allowed, regardless of any rules.

- A. Security groups
- B. Availability Zones
- C. Network ACLs
- D. Geo Redundant Servers

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 408

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What happens if the instance launched by Auto Scaling becomes unhealthy?

- A. Auto Scaling will terminate the instance and launch a new healthy instance.
- B. Auto Scaling will terminate the instance but not launch a new instance.
- C. The instance cannot become unhealthy.
- D. Auto Scaling will notify the user and the user can update the instance.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 409

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has set the Alarm for the CPU utilization > 50%. Due to an internal process, the current CPU utilization will be 80% for 6 hours. How can the user ensure that the CloudWatch alarm does not perform any action?

- A. The user can disable the alarm using the DisableAlarmActions API.
- B. The user can set CloudWatch in a sleep state using the CLI mon-sleep-alarm-action.
- C. The user can pause the alarm from the console.
- D. The user cannot stop the alarm from performing an action unless the alarm is deleted.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 410

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is sending a custom metric to CloudWatch. If the call to the CloudWatch APIs has different dimensions, but the same metric name, how will CloudWatch treat all the requests?

- A. It will treat each unique combination of dimensions as a separate metric.
- B. It will group all the calls into a single call.
- C. It will overwrite the previous dimension data with the new dimension data.
- D. It will reject the request as there cannot be a separate dimension for a single metric.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 411

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has enabled the CloudWatch alarm to estimate the usage charges. If the user disables monitoring of the estimated charges but does not delete the billing alert from the AWS account, what will happen?

- A. The user cannot edit the existing billing alarm.
- B. The data collection on estimated charges is stopped.
- C. It is not possible to disable monitoring of the estimated charges.
- D. AWS will stop sending the billing alerts to the user.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 412

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does enabling a sticky session with ELB do?

- A. Routes all the requests to a single DNS
- B. Ensures that all requests from the user's session are sent to multiple instances
- C. Binds the user session with a specific instance
- D. Provides a single ELB DNS for each IP address

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 413

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following statements is true of an Auto Scaling group?

- A. An Auto Scaling group cannot span multiple regions.
- B. An Auto Scaling group delivers log files within 30 minutes of an API call.
- C. Auto Scaling publishes new log files about every 15 minutes.
- D. An Auto Scaling group cannot be configured to scale automatically.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 414

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following activities is NOT performed by the Auto Scaling policy?

- A. Changing instance types
- B. Scaling up instance counts
- C. Maintaining current instance levels
- D. Scaling down instance counts

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 415

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following services is used to monitor the Amazon Web Services resources?

- A. AWS CloudWatch
- B. AWS Cloudfront
- C. AWS Monitor
- D. AWS EC2

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 416

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is Amazon Import/Export?

- A. A properly configured service role and instance profile
- B. An international shipping division to help you enhance your sales reach
- C. A service that accelerates transferring large amounts of data into and out of AWS using physical storage appliances
- D. A software developed by Amazon to migrate the data from/to your datacenter to AWS

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 417

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the choices below best describes what Auto Scaling is well suited for?

- A. only for applications that experience hourly, daily, or weekly variability in usage.
- B. Both for applications that have stable demand patterns and that experience hourly, daily, or weekly variability in usage.
- C. Both for applications that use frameworks and SDKs to enhance its customer relationship.
- D. only for applications with a stable usage pattern but extremely high workload.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 418

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

True or False: Amazon Route 53 provides highly available and scalable Domain Name System (DNS), domain name registration, and health-checking web services.

- A. False, you can only import an existing domain using Amazon Route 53.
- B. True, however, it only provides .com domains.
- C. FALSE
- D. TRUE

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 419

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following statements is true of Elastic Load Balancing?

- A. It distributes traffic only to instances across different Availability Zones.
- B. It distributes the outgoing traffic across multiple EC2 instances.
- C. It distributes incoming traffic across multiple EC2 instances.
- D. It distributes traffic only to instances across a single Availability Zone.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 420

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You are setting up a VPC and you need to set up a public subnet within that VPC. Which following requirement must be met for this subnet to be considered a public subnet?

- A. Subnet's traffic is not routed to an internet gateway but has its traffic routed to a virtual private gateway.
- B. Subnet's traffic is routed to an internet gateway.
- C. Subnet's traffic is not routed to an internet gateway.
- D. None of these answers can be considered a public subnet.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 421

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following services can receive an alert from CloudWatch?

- A. AWS Elastic Block Store
- B. AWS Relational Database Service
- C. AWS Auto Scaling
- D. AWS Elastic Load Balancing

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 422

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user creates an Auto Scaling group from the Amazon AWS Console and assigned a tag with a key of "environment" and a value of "Prod". Can the user assign tags to instances launched in the Auto Scaling group, to organize and manage them?

- A. Yes, this is possible only if the tags are configured at the launch configuration with a maximum length of 300 characters.
- B. Yes
- C. Yes, this is possible only if the tags are in the same AZ and the tag names are uppercase.
- D. No

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 423

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

How many metrics are supported by CloudWatch for Auto Scaling?

- A. 8 metrics and 1 dimension
- B. 7 metrics and 5 dimension
- C. 5 metrics and 1 dimension
- D. 1 metric and 5 dimensions

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 424

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is aware that a huge download is occurring on his instance. He has already set the Auto Scaling policy to increase the instance count when the network I/O increases beyond a certain limit. How can the user ensure that this temporary event does not result in scaling?

- A. The policy cannot be set on the network I/O
- B. There is no way the user can stop scaling as it is already configured
- C. The network I/O are not affected during data download
- D. He can suspend scaling temporarily

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 425

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send _____ minute metrics to Amazon CloudWatch.

- A. 4
- B. 2
- C. 1
- D. 5

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 426

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following is true of Amazon CloudWatch?

- A. Amazon CloudWatch monitors Amazon Web Services (AWS) resources and the applications that run on AWS in real-time.
- B. Amazon CloudWatch is a web service that gives businesses an easy and cost effective way to distribute content with low latency and high data transfer speeds.
- C. Amazon CloudWatch runs code without provisioning or managing servers.
- D. None of these are true.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 428

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In the AWS Storage Gateway, using the _____, you can cost-effectively and durably archive backup data in Amazon Glacier.

- A. Gateway-virtual tape library (Gateway-VTL)
- B. Gateway-stored volume
- C. Gateway-cached volume
- D. Volume gateway

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 429

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What are the benefits of CloudTrail integration with CloudWatch Logs?

- A. It delivers API activity captured by CloudTrail to an S3 bucket.
- B. It doesn't exist
- C. It delivers SDK activity captured by CloudTrail to a CloudWatch Logs log stream.
- D. It delivers API activity captured by CloudTrail to a CloudWatch Logs log stream.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 430

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Security groups in VPC operate at the _____.

- A. data transport layer level
- B. subnet level
- C. instance level
- D. gateway level

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 431

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Network ACLs are _____.

- A. stateful
- B. stateless
- C. asynchronous
- D. synchronous

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 432

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Is it possible to publish your own metrics to CloudWatch?

- A. Yes, but only if the data is aggregated.
- B. No, it is not possible.
- C. No, metrics are in-built and cannot be defined explicitly.
- D. Yes, it can be done by using the put-metric-data command.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 433

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Can you use CloudWatch to monitor memory and disk utilization usage for your Amazon EC2 Linux instances?

- A. CloudWatch can only measure memory usage.
- B. CloudWatch can only collect memory and disk usage metrics when an instance is running.
- C. It is possible only on Linux EC2 instances using the CloudWatch Monitoring scripts for Linux.
- D. CloudWatch can only measure disk usage.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 434

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An Auto Scaling group is running at the desired capacity of 5 instances and receives a trigger from the Cloudwatch Alarm to increase the capacity by 1. The cool down period is 5 minutes.

Cloudwatch sends another trigger after 2 minutes to decrease the desired capacity by 1. What will be the count of instances at the end of 4 minutes?

- A. 7
- B. 6
- C. 4
- D. 5

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 435

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An instance has enabled basic monitoring only for CloudWatch. What is the minimum time period available for basic monitoring?

- A. 60 seconds
- B. 360 seconds
- C. 300 seconds
- D. 240 seconds

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 436

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following statements describes launch configuration in Auto Scaling?

- A. A launch configuration is a template that an Auto Scaling group uses to launch EC2 instances.
- B. A launch configuration is a template that an Auto Scaling group uses to define the max/minimum of instances.
- C. A launch configuration is a template that an Auto Scaling group uses to schedule the scaling activity.
- D. A launch configuration is a template that an Auto Scaling group uses to define the instance count.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 437

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 438

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Amazon RDS provides Amazon CloudWatch metrics for your DB Instance deployments at no additional charge. You can use the AWS Management Console to view key operational metrics for your DB Instance deployments, including ____.

- A. I/O activity, DB Instance connections, and number of users
- B. DB Engine Version Management
- C. username, I/O activity, and DB Instance connections
- D. compute/memory/storage capacity utilization, I/O activity, and DB Instance connections

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 439

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A custom network ACL that you create _____ until you add rules, and is not associated with a sub-net until you explicitly associate it with one.

- A. blocks only inbound traffic by default
- B. allows outbound traffic by default
- C. allows all inbound and outbound traffic by default
- D. blocks all inbound and outbound traffic by default

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 441

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In AWS Storage Gateway, Gateway-cached volumes allow you to retain _____.

- A. a durable and inexpensive offsite backup that you can recover locally
- B. your primary data locally, and asynchronously back up point-in-time snapshots of this data to Amazon S3
- C. your backup application with online access to virtual tapes
- D. low-latency access to your frequently accessed data

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 442

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following states is not possible for the CloudWatch alarm?

- A. ALERT
- B. ALARM
- C. OK
- D. INSUFFICIENT_DATA

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 443

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is the default maximum number of VPCs allowed per region?

- A. 5
- B. 15
- C. 100
- D. 10

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 444

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

How often is metric data is sent to CloudWatch when detailed monitoring is enabled on an Amazon EC2 instance?

- A. Every 30 seconds
- B. Every 5 minutes
- C. Every 15 minutes
- D. Every minute

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 445

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A route table in VPC can be associated with multiple subnets. However, a subnet can be associated with only _____ route table(s) at a time.

- A. four
- B. two
- C. three
- D. one

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 446

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following statements is NOT true of CloudWatch?

- A. CloudWatch can be accessed using the AWS SDKS.
- B. CloudWatch can be accessed using the AWS console.
- C. CloudWatch can be accessed using CloudWatch API.
- D. CloudWatch can be accessed using the CloudWatch CLI for iOS.

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 447

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following is an incorrect statement about Amazon CloudWatch?

- A. You can use CloudWatch to collect and track metrics, which are the variables you want to measure for your resources and applications.
- B. You can set CloudWatch alarms to send notifications or automatically make changes to the resources you are monitoring, based on rules that you define.
- C. You can control and monitor all Security Groups and their related rules.
- D. You gain system-wide visibility into resource utilization, application performance, and operation-al health.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 448

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following terms is NOT a key CloudWatch concept?

- A. Namespaces
- B. Units
- C. Time Stamps
- D. Indexes

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 449

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Network ACLs in a VPC operate at the _____.

- A. TCP level
- B. instance level
- C. subnet level
- D. gateway level

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 450

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Your VPC automatically comes with a modifiable default network ACL, which by default _____.

- A. blocks outbound traffic
- B. allows only inbound traffic
- C. allows all inbound and outbound traffic
- D. blocks all inbound and outbound traffic

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 451

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is a placement group in Amazon EC2?

- A. It is a logical grouping of EC2 instances within a single Availability Zone.
- B. It the edge location of your web content.
- C. It is a group used to span multiple Availability Zones.
- D. It is the AWS region where you run the EC2 instance of your web content.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 452

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In which screen does a user select the Availability Zones while configuring Auto Scaling?

- A. Auto Scaling Group Creation
- B. Auto Scaling Instance Creation
- C. Auto Scaling Launch config Creation
- D. Auto Scaling Policy Creation

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 453

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the CloudWatch services mentioned below is NOT a part of the AWS free tier?

- A. 10 alarms/month
- B. 1 million API request/month
- C. 10 metrics/month
- D. 15 detailed monitoring metrics

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 454

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In the context of sending metrics to CloudWatch using Amazon Kinesis, which of the following statements best describes the metric "PutRecord.Latency"?

- A. It is the time taken per PutRecord operation, measured over the specified time period.
- B. It is the number of successful records in a PutRecords operation per Amazon Kinesis stream, measured over the specified time period.
- C. It is the time taken per PutRecords operation to calculate the statistics of the PutRecords operations.
- D. It is the number of successful PutRecord operations per Amazon Kinesis stream, measured over the specified time period.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 455

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Can a user depict CloudWatch metrics such as CPU utilization in % and Network I/O in bytes on a single graph?

- A. No, a user cannot graph two separate metrics on the same graph.
- B. Yes, a user can graph several metrics over time on a single graph.
- C. No, a user cannot plot several metrics on a single graph since the units are different.
- D. Yes, a user can graph multiple metrics on the same graph provided they are of the same instance in the same AZ.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 456

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following comes before Auto Scaling group creation?

- A. Creating the Auto Scaling launch config
- B. Creating the Auto Scaling policy
- C. Creating the Auto Scaling tags
- D. Creating the Auto Scaling instance

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 457

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A placement group in Amazon EC2 can

- A. place high memory instances in one logical group.
- B. logically name and tag different tiers of the system (DB, application, business logic etc).
- C. isolate any instance-type physically so that groups access local resources.
- D. reduce network latency and increase network throughput

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 458

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following statements is true about Auto Scaling?

- A. You can only delete your Auto Scaling group but not your Auto Scaling setup.
- B. If the Auto Scaling infrastructure is being deleted, it is not mandatory to delete the launch configuration.
- C. You can only delete your Auto Scaling set up but not your Auto Scaling group.
- D. If the Auto Scaling infrastructure is being deleted, it is mandatory to delete the launch configuration.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 459

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

If you specify only the general endpoint (autoscaling.amazonaws.com), Auto Scaling directs your request to the:

- A. us-west-2 endpoint.
- B. eu-central-1.
- C. eu-west-1 endpoint.
- D. us-east-1 endpoint.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 460

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has configured ELB with Auto Scaling. The user temporarily suspended the Auto Scaling terminate process. What might the Availability Zone Rebalancing process (AZRebalance) consequently cause during this period?

- A. Auto Scaling will keep launching instances in all AZs until the maximum instance number is reached.
- B. AZ Rebalancing might now allow Auto Scaling to launch or terminate any instances.
- C. AZ Rebalancing might allow the number instances in an Availability Zone to remain higher than the maximum size
- D. It is not possible to suspend the terminate process while keeping the launch active.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 461

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What is Amazon CloudFront?

- A. A global Content Delivery Network
- B. An encrypted endpoint to upload files to the Cloud
- C. A web service to schedule regular data movement
- D. A development front-end to Amazon Web Services

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 462

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You can create a CloudWatch alarm that watches a single metric. The alarm performs one or more actions based on the value of the metric relative to a threshold over a number of time periods. Which of the following states is possible for the CloudWatch alarm?

- A. OK
- B. ALERT
- C. THRESHOLD
- D. ERROR

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 463

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has enabled instance protection for his Auto Scaling group that has spot instances. If Auto Scaling wants to terminate an instance in this Auto Scaling group due to a CloudWatch trigger unre-lated to bid price, what will happen?

- A. Auto Scaling will notify the user for the next action
- B. Auto Scaling will remove the instance from the Auto Scaling Group
- C. Auto Scaling overwrites the instance termination attribute and terminates the instances
- D. The EC2 instance will not be terminated since instance protection from scale-in is enabled.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 464

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In a hardware security module (HSM), what is the function of a Transparent Data Encryption (TDE)?

- A. To reduce the risk of confidential data theft
- B. To decrease latency
- C. To store SSL certificates
- D. To provide backup

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 465

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In IAM, a policy has to include the information about who (user) is allowed to access the resource, known as the ____.

- A. permission
- B. role
- C. license
- D. principal

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 466

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Amazon S3 provides a number of security features for protection of data at rest, which you can use or not, depending on your threat profile. What feature of S3 allows you to create and manage your own encryption keys for sending data?

- A. Client-side Encryption
- B. Network traffic protection
- C. Data integrity compromise
- D. Server-side Encryption

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 468

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

AWS Cloud Hardware Security Modules (HSMs) are designed to ____.

- A. store your AWS keys safely
- B. provide another level of login security specifically for LDAP
- C. allow AWS to audit your infrastructure
- D. securely store cryptographic key material and use the key material without exposing it outside the cryptographic boundary of the appliance

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 469

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following statements is true of IAM?

- A. If you are configuring MFA for a user who will use a smartphone to generate an OTP, you must have the smartphone available in order to finish the wizard.
- B. If you are configuring MFA for a user who will use a smartphone to generate an OTP, the smartphone is not required in order to finish the wizard.
- C. If you are configuring MFA for a user who will use a smartphone to generate an OTP, you can finish the wizard on any device and later use the smartphone for authentication.
- D. None of these are correct.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 470

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Could you use IAM to grant access to Amazon DynamoDB resources and API actions?

- A. In DynamoDB there is no need to grant access
- B. Depended to the type of access
- C. No
- D. Yes

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 471

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is planning to schedule a backup for an existing EBS volume. The user wants the backup to be created through snapshot, and for it to be encrypted. How can the user achieve data encryption with a snapshot?

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 472

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does the Server-side encryption provide in Amazon S3?

- A. Server-side encryption doesn't exist for Amazon S3, but only for Amazon EC2.
- B. Server-side encryption protects data at rest using Amazon S3-managed encryption keys (SSE-S3).
- C. Server-side encryption provides an encrypted virtual disk in the cloud.
- D. Server-side encryption allows to upload files using an SSL endpoint for a secure transfer.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 473

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The Statement element, of an AWS IAM policy, contains an array of individual statements. Each individual statement is a(n) _____ block enclosed in braces { }.

- A. JSON
- B. AJAX
- C. JavaScript
- D. jQuery

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 474

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: Security groups and network access control lists (ACLs). You start to look into security groups first. Which statement below is incorrect in relation to security groups?

- A. Are stateful: Return traffic is automatically allowed, regardless of any rules.
- B. Support addition of individual allow and deny rules in both inbound and outbound.
- C. Security Groups can be added or removed from EC2 instances in a VPC at any time.
- D. Evaluate all rules before deciding whether to allow traffic.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 475

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon IAM stand for?

- A. Amazon Identity and Authentication Mechanism
- B. Amazon Integrated Access Management
- C. Amazon Identity and Access Management
- D. None of these

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 476

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Can you use the AWS Identity and Access Management (IAM) to assign permissions determining who can manage or modify RDS resources?

- A. No, AWS IAM is used only to assign IDs to AWS users.
- B. No, this permission cannot be assigned by AWS IAM.
- C. Yes, you can.
- D. No, AWS IAM is used only to assign activities.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 477

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You have been asked to design a layered security solution for protecting your organization's network infrastructure. You research several options and decide to deploy a network-level security control appliance, inline, where traffic is intercepted and analyzed prior to being forwarded to its final destination, such as an application server. Which of the following is NOT considered an inline threat protection technology?

- A. Intrusion prevention systems
- B. Third-party firewall devices installed on Amazon EC2 instances
- C. Data loss management gateways
- D. Augmented security groups with Network ACLs

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 478

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Is it possible to protect the connections between your application servers and your MySQL instances using SSL encryption?

- A. Yes, it is possible but only in certain regions.
- B. Yes
- C. No
- D. Yes, it is possible but only in VPC.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 479

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You need to determine what encryption operations were taken with which key in AWS KMS to either encrypt or decrypt data in the AWS CodeCommit repository. Which of the following actions will best help you accomplish this?

- A. Searching for the AWS CodeCommit repository ID in AWS CloudTrail logs
- B. Searching for the encryption key ID in AWS CloudTrail logs
- C. Searching for the AWS CodeCommit repository ID in AWS CloudWatch
- D. Searching for the encryption key ID in AWS CloudWatch

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 480

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The AWS Key Management Service (AWS KMS) is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data. AWS KMS is integrated with other AWS services including Amazon EBS, Amazon S3, Amazon Redshift, Elastic Transcoder, Amazon WorkMail, and Amazon RDS to make it simple to encrypt your data with encryption keys that you manage. AWS KMS is also integrated with AWS CloudTrail to provide you with key usage logs to help meet your regulatory and compliance needs. Which of the following types of cryptography keys is supported by AWS KMS currently?

- A. Private ephemeral key agreement cryptography
- B. Symmetric and asymmetric random number generation key cryptography
- C. Asymmetric key cryptography and symmetric key cryptography
- D. Only symmetric key cryptography

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 481

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Your customers are concerned about the security of their sensitive data and their inquiry asks about what happens to old storage devices on AWS. What would be the best answer to this question?

- A. AWS uses a 3rd party security organization to destroy data as part of the decommissioning process.
- B. AWS uses the techniques detailed in DoD 5220.22-M to destroy data as part of the decommissioning process.
- C. AWS reformats the disks and uses them again.
- D. AWS uses their own proprietary software to destroy data as part of the decommissioning process.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 482

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In AWS Identity and Access Management (IAM), you can make use of the _____ APIs to grant users temporary access to your resources.

- A. AWS Security Transport Service (STS)
- B. AWS Security Tree Service (STS)
- C. AWS Security Task Service (STS)
- D. AWS Security Token Service (STS)

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 483

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An IAM user has two conflicting policies as part of two separate groups. One policy allows him to access an S3 bucket, while another policy denies him the access. Can the user access that bucket?

- A. Yes, always
- B. No
- C. Yes, provided he accesses with the group which has S3 access
- D. Yes, but just read only access of the bucket

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 484

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization has launched 5 instances: 2 for production and 3 for testing. The organization wants a particular group of IAM users to access only the test instances and not the production ones. They want to deploy the instances in various locations based on the factors that will change from time to time, especially in the test group. They expect instances will often need to be churned, i.e. deleted and replaced, especially in the testing group. This means the five instances they have created now will soon be replaced by a different set of five instances. The members of each group, production and testing, will not change in the foreseeable future. Given the situation, what choice below is the most efficient and time-saving strategy to define the IAM policy?

- A. By creating an IAM policy with a condition that allows access to only small instances
- B. By defining the IAM policy that allows access based on the instance ID
- C. By launching the test and production instances in separate regions and allowing region wise access to the group
- D. By defining the tags on the test and production team members IAM user IDs, and adding a condition to the IAM policy that allows access to specific tags

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 485

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

For IAM user, a virtual Multi-Factor Authentication (MFA) device uses an application that generates ____-digit authentication codes that are compatible with the time-based one-time password (TOTP) standard.

- A. three
- B. four
- C. six
- D. five

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 486

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The ____ IAM policy element describes the specific action or actions that will be allowed or de-nied.

- A. Principal
- B. Action
- C. Vendor
- D. Not Principal

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 487

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company wants to review the security requirements of Glacier. Which of the below mentioned statements is true with respect to the AWS Glacier data security?

- A. The user can set the serverside encryption flag to encrypt the data stored on Glacier.
- B. All data stored on Glacier is protected with AES-256 server-side encryption.
- C. All data stored on Glacier is protected with AES-128 server-side encryption.
- D. The data stored on Glacier is not encrypted by default.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 488

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user has configured two security groups which allow traffic as given below:

1: SecGrp1:

Inbound on port 80 for 0.0.0.0/0

Inbound on port 22 for 0.0.0.0/0

2: SecGrp2:

Inbound on port 22 for 10.10.10.1/32

If both the security groups are associated with the same instance, which of the below mentioned statements is true?

- A. It is not possible to have more than one security group assigned to a single instance
- B. It allows inbound traffic for everyone on both ports 22 and 80
- C. It is not possible to create the security group with conflicting rules. AWS will reject the request
- D. It allows inbound traffic on port 22 for IP 10.10.10.1 and for everyone else on port 80

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 489

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Is it possible to create an S3 bucket accessible only by a certain IAM user using policies in a Cloud-Formation template?

- A. Yes, all these resources can be created using a CloudFormation template
- B. S3 is not supported by CloudFormation.
- C. No, you can only create the S3 bucket but not the IAM user.
- D. No, in the same template you can only create the S3 bucket and the relative policy.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 490

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In Amazon CloudFront, if you have chosen On for Logging, the access logs are stored in _____.

- A. Amazon S3 bucket.
- B. Amazon EBS.
- C. Amazon Edge locations.
- D. Amazon EC2 instance.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 491

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Does Amazon RDS support SSL encryption for SQL Server DB Instances?

- A. Yes, for all supported SQL Server editions
- B. No
- C. Yes, but only when the instances are in a single region
- D. No, encryption using SSL is supported only in the GovCloud.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 492

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Amazon Cognito supports web identity federation through _____.

- A. custom sign-in code or own user identities
- B. Facebook, Google, and Amazon
- C. a configuration check for rules that deny access to specific ports
- D. an AWS user group

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 493

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

- A. The user should create an IAM user with permissions to access DynamoDB and use its credentials within the application for connecting to DynamoDB
- B. The user should create an IAM user with DynamoDB and EC2 permissions. Attach the user with the application so that it does not use the root account credentials
- C. The user should attach an IAM role to the EC2 instance with necessary permissions for making API calls to DynamoDB.
- D. The user should create an IAM role with EC2 permissions to deploy the application

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 494

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A user is trying to create a list of IAM users with the AWS console. When the IAM users are created which of the below mentioned credentials will be enabled by default for the user?

- A. IAM X.509 certificates
- B. Nothing. Everything is disabled by default
- C. IAM passwords
- D. IAM access key and secret access key

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 495

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You are setting up security groups for both incoming traffic and outgoing traffic in your VPC network on the AWS CLI. Which of the following AWS CLI commands would you use for adding one or more incoming traffic rules to a security group?

- A. authorize-security-group-egress
- B. authorize-security-group-ingress
- C. Grant-EC2SecurityGroupOutgress
- D. Get-EC2SecurityGroup

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 496

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The IAM entity "AWS Account" is similar to:

- A. The Unix concept of root or superuser
- B. The Unix concept of a non privilege user
- C. The Unix concept of guest user
- D. The primary billing entity

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 497

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

AWS KMS (Key Management Service) uses symmetric key cryptography to perform encryption and decryption. Symmetric key cryptography uses the same algorithm and key to both encrypt and de-crypt digital data. The unencrypted data is typically called plaintext whether it is text or not, and the encrypted data is typically called _____.

- A. ciphertext
- B. symtext
- C. encryptext
- D. cryptext

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 498

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Bob is an IAM user who has access to the EC2 services. Admin is an IAM user who has access to all the AWS services including IAM. Can Bob change his own password?

- A. No, the IAM user can never change the password
- B. Yes, only from AWS CLI
- C. Yes, only from the AWS console
- D. Yes, provided Admin has given Bob access to change his own password

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 499

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

ABC has three AWS accounts. They have created separate IAM users within each account.

ABC wants a single IAM login URL such as `https://abc.signin.aws.amazon.com/console/` for use by IAM users in all three accounts.

How can this be achieved?

- A. Merge all the accounts with consolidated billing
- B. Create the S3 bucket with an alias name and use the redirect rule to forward requests to various accounts
- C. Create the same account alias with each account ID
- D. It is not possible to have the same IAM account login URL for separate AWS accounts

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 500

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the following Identity and Access Management (IAM) policy keys of AWS Direct Connect is used for date/time conditions?

- A. aws:CurrentTime
- B. aws:UserAgent
- C. aws:SourceIp
- D. aws:SecureTransport

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 501

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In the context of AWS Security Best Practices for RDS, if you require encryption or data integrity authentication of data at rest for compliance or other purposes, you can add protection at the _____ using SQL cryptographic functions.

- A. physical layer
- B. security layer
- C. application layer
- D. data-link layer

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 502

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A root AWS account owner has created three IAM users: Bob, John and Michael. Michael is the IAM administrator. Bob and John are not the super users, but users with some pre-defined policies. John does not have access to modify his password. Thus, he asks Bob to change his password. How can Bob change John's password?

- A. This statement is false. Only Michael can change the password for John
- B. This is possible if Michael can add Bob to a group which has permissions to modify the IAM passwords
- C. It is not possible for John to modify his password
- D. Provided Bob is the manager of John

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 504

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

What does Amazon IAM provide?

- A. A mechanism to authorize Internet Access Modularity (IAM)
- B. A mechanism to authenticate users when accessing Amazon Web Services
- C. A mechanism to integrate on-premises authentication protocols with the Cloud
- D. None of the above

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 505

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An IAM group is a:

- A. group of EC2 machines that gain the permissions specified in the group.
- B. collection of IAM users.
- C. guide for IAM users.
- D. collection of AWS accounts.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 506

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A group in IAM can contain many users. Can a user belong to multiple groups?

- A. Yes, a user can be a member of up to 150 groups.
- B. Yes, a user can be a member of up to 50 groups.
- C. Yes, a user can be a member of up to 100 groups.
- D. Yes, a user can be a member of up to 10 groups.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 507

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Fill in the blanks: One of the basic characteristics of security groups for your VPC is that you

----- .

- A. can specify allow rules as well as deny rules
- B. can neither specify allow rules nor deny rules
- C. can specify allow rules, but not deny rules
- D. can specify deny rules, but not allow rules

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 508

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You can configure Amazon CloudFront to deliver access logs per _____ to an Amazon S3 bucket of your choice.

- A. Edge location
- B. Distribution
- C. Geo restriction
- D. Request

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 509

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

ABC (with AWS account ID 111122223333) has created 50 IAM users for its organization's employees. What will be the AWS console URL for these associates?

- A. <https://signin.aws.amazon.com/console/111122223333/>
- B. <https://111122223333.signin.aws.amazon.com/console/>
- C. <https://signin.aws.amazon.com/111122223333/console/>
- D. <https://signin.aws.amazon.com/console/>

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 510

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

AWS IAM permissions can be assigned in two ways:

- A. as role-based or as resource-based.
- B. as identity-based or as resource-based.
- C. as security group-based or as key-based.
- D. as user-based or as key-based.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 511

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Can you change the security groups associated with the primary network interface (eth0) of an EC2 instance running inside a VPC?

- A. Yes
- B. Only if the instance is stopped
- C. Only when the instance is launched
- D. No

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 512

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Amazon Relational Database Service integrates with _____, a service that lets your organization create users and groups under your organization's AWS account and assign unique security credentials to each user.

- A. Amazon RDS tags
- B. AWS IAM
- C. AWS Lambda
- D. Amazon EMR

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 513

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The information within an IAM policy is described through a series of _____.

- A. elements
- B. macros
- C. classes
- D. namespaces

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 514

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In Amazon VPC, the _____ encryption function is used to ensure privacy among both IKE and IPsec Security Associations.

- A. AES 192-bit
- B. AES 256-bit
- C. SHA 180-bit
- D. SHA 2-bit

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 515

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In IAM, can you attach more than one inline policy to a particular entity such a user, role, or group?

- A. No
- B. Yes
- C. Yes, you can but only if you attach the policy within a VPC.
- D. Yes, you can but only if you attach the policy within the GovCloud.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 516

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A customer enquires about whether all his data is secure on AWS, and is especially concerned about Elastic Map Reduce (EMR). You need to inform him of some of the security features in place for AWS. Which of the below statements is incorrect regarding EMR or S3?

- A. Every packet sent in the AWS network uses Internet Protocol Security (IPsec).
- B. Amazon S3 provides authentication mechanisms to ensure that stored data is secured against un-authorized access.
- C. Customers may encrypt the input data before they upload it to Amazon S3.
- D. Amazon EMR customers can choose to send data to Amazon S3 using the HTTPS protocol for secure transmission.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 517

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

If an IAM policy has multiple conditions, or if a condition has multiple keys, its boolean outcome will be calculated using a logical _____ operation.

- A. NAND
- B. OR
- C. AND
- D. None of these

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 518

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

You have set up an IAM policy for your users to access Elastic Load Balancers and you know that an IAM policy is a JSON document that consists of one or more statements. Which of the following elements is not a part of the statement in an IAM policy document?

- A. Action
- B. Resource
- C. Effect
- D. Key

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 519

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

In AWS Identity and Access Management, roles can be used by an external user authenticated by an external identity provider (IdP) service that is compatible with _____.

- A. BNML (Business Narrative Markup Language)
- B. CFML (ColdFusion Markup Language)
- C. SAML 2.0 (Security Assertion Markup Language 2.0)
- D. BPML (Business Process Modeling Language)

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 520

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Which of the below mentioned options is not a best practice to securely manage the AWS access credentials?

- A. Keep rotating your secure access credentials at regular intervals
- B. Create individual IAM users
- C. Create strong access key and secret access key and attach to the root account
- D. Enable MFA for privileged users

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 521

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The amount of data a company must back up has been increasing, and storage space is quickly running out. There is no budget to purchase new backup software that is capable of backing up data directly to the cloud.

What is the MOST cost-effective way to make storage available to the company's legacy backup system?

- A. Launch an Amazon EC2 instance, add large Amazon EBS volumes, and connect using VPN
- B. Ship backup tapes to AWS for storage in secure AWS Availability Zones
- C. Use AWS Snowball on a weekly basis to transfer data to Amazon Glacier
- D. Use AWS Storage Gateway to present a VTL using iSCSI to the legacy application

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 522

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The SysOps Administrator must integrate an existing on-premises asymmetrical key management system into an AWS services platform. How can the Administrator meet this requirement?

- A. Implement AWS KMS and integrate with the existing on-premises asymmetrical key management system
- B. Implement AWS CloudHSM and integrate it with the existing key management infrastructure
- C. Deploy an Amazon EC2 instance and choose an AMI from an AWS partner in the AWS Marketplace
- D. Create a master key in AWS KMS, and export that key to the existing on-premises asymmetrical key management system

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 523

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A Systems Administrator is planning to deploy multiple EC2 instances within two separate Availability Zones in the same AWS Region. The instances cannot be exposed to the Internet, but must be able to exchange traffic between one another. The data does not need to be encrypted.

What solution meets these requirements while maintaining the lowest cost?

- A. Create two private subnets within the same VPC. Communicate between instances using their private IP addresses
- B. Create 2 public subnets within the same VPC. Communicate between instances using their public IP addresses
- C. Create 2 separate VPCs, one for each Availability Zone. Create a private subnet within each VPC. Create a static route table pointing the destination CIDR to the other VPC
- D. Create 2 separate VPCs, one for each Availability Zone and create a public subnet in each. Deploy a VPN appliance within each VPC and establish a VPN tunnel between them. Communicate between instances by routing traffic through the VPN appliances

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 524

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company website hosts patches for software that is sold globally. The website runs in AWS and performs well until a large software patch is released. The flood of downloads puts a strain on the web servers and leads to a poor customer experience.

What can the SysOps Administrator propose to enhance customer experience, create a more available web platform, and keep costs low?

- A. Use an Amazon CloudFront distribution to cache static content, including software patches
- B. Increase the size of the NAT instance to improve throughput
- C. Scale out of web servers in advance of patch releases to reduce Auto Scaling delays
- D. Move the content to IO1 and provision additional IOPS to the volume that contains the software patches

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 525

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization has developed a new memory-intensive application that is deployed to a large Amazon EC2 Linux fleet. There is concern about potential memory exhaustion, so the Development team wants to monitor memory usage by using Amazon CloudWatch.

What is the MOST efficient way to accomplish this goal?

- A. Deploy the solution to memory-optimized EC2 instances, and use the CloudWatch MemoryUtilization metric
- B. Enable the Memory Monitoring option by using AWS Config
- C. Install the AWS Systems Manager agent on the applicable EC2 instances to monitor memory
- D. Monitor memory by using a script within the instance, and send it to CloudWatch as a custom metric

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 526

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator is running Amazon EC2 instances in multiple AWS Regions. The Administrator wants to aggregate the CPU utilization for all instances onto an Amazon CloudWatch dashboard. Each region should be present on the dashboard and represented by a single graph that contains the CPU utilization for all instances in that region.

How can the Administrator meet these requirements?

- A. Create a cross-region dashboard using AWS Lambda and distribute it to all regions
- B. Create a custom CloudWatch dashboard and add a widget for each region in the AWS Management Console
- C. Enable cross-region dashboards under the CloudWatch section of the AWS Management Console
- D. Switch from basic monitoring to detailed monitoring on all instances

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 527

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A mobile application must allow users to securely access their own content stored in a shared Amazon S3 bucket.

Which AWS services should be used to enable this access? (Choose two.)

- A. AWS Directory Service
- B. AWS Shield
- C. IAM roles
- D. Amazon Cognito
- E. AWS Organizations

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 528

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A Development team has an application stack consisting of many OS dependencies and language runtime dependencies. When deploying the application to production, the most important factor is how quickly the instance is operational.

What deployment methodology should be used to update the running environments to meet the requirement?

- A. Use fully baked AMIs (λgolden imagesλ) created after each successful build, creating a new Auto Scaling group, and blue/green deployments with rollbacks.
- B. Use user-data scripts to configure the instance correctly on boot by installing all dependencies when needed.
- C. Use an AWS Lambda function to only update the application locally on each instance, then re-attach it to the load balancer when the process complete.
- D. Use AWS OpsWorks scripts to execute on reboot of each instance to install all known dependencies, then re-attach the instances to the load balancer.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 529

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A web-based application is running in AWS. The application is using a MySQL Amazon RDS database instance for persistence. The application stores transactional data and is read-heavy. The RDS instance gets busy during the peak usage, which shows the overall application response times.

The SysOps Administrator is asked to improve the read queries performance using a scalable solution.

Which options will meet these requirements? (Choose two.)

- A. Scale up the RDS instance to a larger instance size
- B. Enable the RDS database Multi-AZ option
- C. Create a read replica of the RDS instance
- D. Use Amazon DynamoDB instead of RDS
- E. Use Amazon ElastiCache to cache read queries

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 530

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A Content Processing team has notified a SysOps Administrator that their content is sometimes taking a long time to process, whereas other times it processes quickly. The Content Processing submits messages to an Amazon Simple Queue Service (Amazon SQS) queue, which details the files that need to be processed.

An Amazon EC2 instance polls the queue to determine which file to process next.

How could the Administrator maintain a fast but cost-effective processing time?

- A. Attach an Auto Scaling policy to the Amazon SQS queue to increase the number of EC2 instances based on the depth of the SQS queue
- B. Create an Auto Scaling policy to increase the number of EC2 instances polling the queue and a CloudWatch alarm to scale based on MaxVisibility Timeout
- C. Attach an Auto Scaling policy to the SQS queue to scale instances based on the depth of the dead-letter queue
- D. Create an Auto Scaling policy to increase the number of EC2 instances polling the queue and a CloudWatch alarm to scale based on ApproximateNumberOfMessagesVisible

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 531

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator receives reports of an Auto Scaling group failing to scale when the nodes running Amazon Linux in the cluster are constrained by high memory utilization.

What should the Administrator do to enable scaling to better adapt to the high memory utilization?

- A. Create a custom script that pipes memory utilization to Amazon S3, then, scale with an AWS Lambda-powered event
- B. Install the Amazon CloudWatch memory monitoring scripts, and create a custom metric based on the script's results
- C. Increase the minimum size of the cluster to meet memory and application load demands
- D. Deploy an Application Load Balancer to more evenly distribute traffic among nodes

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 532

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator has received a request from the Compliance Department to enforce encryption at rest of all new objects uploaded to the corp-compliance bucket.

How can the Administrator enforce encryption on all objects uploaded to the bucket?

A. Enable Amazon S3 default encryption on the bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:PutObject"
      ],
      "Effect": "Deny",
      "Resource": "arn:aws:s3:::corp-compliance/",
      "Condition": {
        "Bool": {
          "aws:SecureTransport": "false"
        }
      },
      "Principal": "*"
    }
  ]
}
```

B. Add the following policy statement to the bucket:

C. Add the following policy statement to the IAM user permissions policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::corp-compliance/",
      "Condition": {
        "Bool": {
          "aws:SecureTransport": "false"
        }
      }
    }
  ]
}
```

D. Generate a presigned URL for the Amazon S3 PUT operation with server-side encryption flag set, and send the URL to the user.

Show Suggested Answer

Exam question from Amazon's AWS-SysOps

Question #: 533

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An errant process is known to use an entire processor and run at 100%. A SysOps Administrator wants to automate restarting the instance once the problem occurs for more than 2 minutes.

How can this be accomplished?

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

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 534

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator needs to report on Amazon EC2 instance cost by both project and environment (production, staging, development). Which action would impact the operations team the LEAST?

- A. For each project and environment, create a new AWS account and link them to the master payer for unified management and billing
- B. Use AWS Organizations to create a new organization for each project, then for each environment use a separate linked AWS account
- C. Implement cost allocation tagging in the Billing and Cost Management console to implement tags to identify resources by project and environment
- D. Add the project and environment information to the instance metadata so that the values can be queried and rolled up into reports

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 535

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A web application's performance has been degrading. Historically, the application has had highly-variable workloads, but lately, there has been a steady growth in traffic as the result of a new product launch. After reviewing several Amazon CloudWatch metrics, it is discovered that over the last two weeks the balance of CPU credits has dropped to zero several times.

Which solutions will improve performance? (Choose two.)

- A. Begin using the T2 instance type
- B. Purchase more CPU credits for the existing instance
- C. Increase the size of the current instance type
- D. Configure a CloudWatch alarm on the CPU credits metric

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 536

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An Amazon EC2 instance is in a private subnet. To SSH to the instance, it is required to use a bastion host that has an IP address of 10.0.0.5. SSH logs on the EC2 instance in the private subnet show that connections are being made over SSH from several other IP addresses. The EC2 instance currently has the following inbound security group rules applied:

Protocol: TCP -

Port: 22 -

Source: 10.0.0.5/32 -

Protocol: TCP -

Port: 22 -

Source: sg-xxxxxxx -

Protocol: TCP -

Port: 389 -

Source: 0.0.0.0/0 -

What is the MOST likely reason that another IP addresses is able to SSH to the EC2 instance?

- A. The rule with 0.0.0.0/0 means SSH is open for any client to connect
- B. The rule with /32 is not limiting to a single IP address
- C. Any instance belonging to sg-xxxxxxx is allowed to connect
- D. There is an outbound rule allowing SSH traffic

Show Suggested Answer

Exam question from Amazon's AWS-SysOps

Question #: 537

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An AWS CloudFormation template creates an Amazon RDS instance. This template is used to build up development environments as needed and then delete the stack when the environment is no longer required. The RDS-persisted data must be retained for further use, even after the CloudFormation stack is deleted.

How can this be achieved in a reliable and efficient way?

- A. Write a script to continue backing up the RDS instance every five minutes
- B. Create an AWS Lambda function to take a snapshot of the RDS instance, and manually execute the function before deleting the stack
- C. Use the Snapshot Deletion Policy in the CloudFormation template definition of the RDS instance
- D. Create a new CloudFormation template to perform backups of the RDS instance, and run this template before deleting the stack

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 538

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company's IT Security team is performing an audit of the AWS environment to determine which servers need to be patched and where additional security controls need to be added.

The company is responsible for which of the following? (Choose two.)

- A. Patching the OS on Amazon RDS instances
- B. Patching the OS on Amazon EC2 instances
- C. Enabling server-side encryption with Amazon S3-Managed Keys (SSE-S3) on S3 objects
- D. Patching the database engine on RDS instances
- E. Patching PHP in an AWS Elastic Beanstalk managed EC2 application

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 539

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The InfoSec team has asked the SysOps Administrator to perform some hardening on the company Amazon RDS database instances. Based on this requirement, what actions should be recommended for the start of the security review? (Choose two.)

- A. Use Amazon Inspector to present a detailed report of security vulnerabilities across the RDS database fleet
- B. Review the security group's inbound access rules for least privilege
- C. Export AWS CloudTrail entries detailing all SSH activity on the RDS instances
- D. Use the cat command to enumerate the allowed SSH keys in ~/.ssh on each RDS instance
- E. Report on the Parameter Group settings and ensure that encrypted connections are enforced

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 540

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A Big Data consulting company wants to separate its customers' workloads for billing and security reasons. The company would like to maintain billing and security controls on these workloads.

According to best practices, how can the workloads be separated if no shared resources are needed?

- A. Require each customer to create their own account. Contact AWS Support to receive a consolidated bill.
- B. Create customer accounts within AWS Organizations specifying consolidated billing features.
- C. Create a separate VPC for each customer. Use security groups to isolate traffic.
- D. Dedicate an AWS Region to each customer. Ensure that each entry in Amazon Route 53 is unique.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 541

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization stores files on Amazon S3. Employees download the files, edit them with the same file name to the same folder on Amazon S3. Occasionally the files are unintentionally modified or deleted.

What is the MOST cost-effective way to ensure that these files can be recovered to their correct state?

- A. Enable cross-region replication on the Amazon S3 bucket
- B. Enable versioning on the Amazon S3 bucket
- C. Use Lifecycle Management to move the files to Amazon Glacier
- D. Copy the edited files to Amazon Elastic File System

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 542

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company has a web application that runs both on-premises and on Amazon EC2 instances. Over time, both the on-premises servers and EC2 instances begin crashing. A SysOps Administrator suspects a memory leak in the application and wants a unified method to monitor memory utilization over time. How can the Administrator track both the EC2 memory utilization and on-premises server memory utilization over time?

- A. Write a script or use a third-party application to report memory utilization for both EC2 instances and on-premises servers.
- B. Use Amazon CloudWatch agent for both Amazon EC2 instances and on-premises servers to report MemoryUtilization metrics to CloudWatch and set a CloudWatch alarm for notifications.
- C. Use CloudWatch agent for Amazon EC2 instances to report memory utilization to CloudWatch, and set CloudWatch alarms for notifications. Use a third-party application for the on-premises servers.
- D. Configure a load balancer to route traffic to both on-premises servers and EC2 instances, then use CloudWatch as the unified view of the metrics for the load balancer.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 543

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Website users report that an application's pages are loading slowly at the beginning of the workday. The application runs on Amazon EC2 instances, and data is stored in an Amazon RDS database. The SysOps Administrator suspects the issue is related to high CPU usage on a component of this application.

How can the Administrator find out which component is causing the performance bottleneck?

- A. Use AWS CloudTrail to review the resource usage history for each component.
- B. Use Amazon CloudWatch metrics to examine the resource usage of each component.
- C. Use Amazon Inspector to view the resource usage details for each component.
- D. Use Amazon CloudWatch Events to examine the high usage events for each component.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 544

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator has an AWS Direct Connect connection in place in region us-east-1, between an AWS account and a data center. The Administrator is now required to connect the data center to a VPC in another AWS Region, us-west-2, which must have consistent network performance and low-latency.

What is the MOST efficient and quickest way to establish this connectivity?

- A. Create an AWS VPN CloudHub architecture, and use software VPN to connect to the VPC in region us-west-2.
- B. Create a new Direct Connect connection between the data center and region us-west-2.
- C. Create a VPC peering connection between the VPC in region us-east-1 and us-west-2, and access the VPC in us-west-2 from the data center.
- D. Use Direct Connect gateway with the existing Direct Connect connection to connect to the Virtual Private Gateway of the VPC in region us-west-2.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 545

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A new application is being tested for deployment on an Amazon EC2 instance that requires greater IOPS than currently provided by the single 4TB General Purpose SSD (gp2) volume.

Which actions should be taken to provide additional Amazon EBS IOPS for the application? (Choose two.)

- A. Increase the size of the General Purpose (gp2) volume
- B. Use RAID 0 to distribute I/O across multiple volumes
- C. Migrate to a Provisioned IOPS SSD (io1) volume
- D. Enable MAX I/O performance mode on the General Purpose (gp2) volume
- E. Use RAID 1 to distribute I/O across multiple volumes

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 546

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A web service runs on Amazon EC2 instances behind an Elastic Load Balancing (ELB) load balancer. External clients must whitelist specific public IP addresses in their firewalls to access the service.

What load balancer or ELB feature should be used for this application?

- A. Network Load Balancer
- B. Application Load Balancer
- C. Classic Load Balancer
- D. Load balancer target groups

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 547

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

While creating the wait condition resource in AWS CloudFormation, a SysOps Administrator receives the error `received 0 signals out of the 1 expected from the EC2 instance`.

What steps should be taken to troubleshoot this issue? (Choose two.)

- A. Confirm from the cfn logs that the cfn-signal command was successfully run on the instance.
- B. Try to re-create the stack with a different IAM user.
- C. Check that the instance has a route to the Internet through a NAT device.
- D. Update the AWS CloudFormation stack service role to have iam:PassRole permission.
- E. Delete the existing stack and attempt to create a new one.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 548

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An existing, deployed solution uses Amazon EC2 instances with Amazon EBS General Purpose SSD volumes, an Amazon RDS PostgreSQL database, an Amazon EFS file system, and static objects stored in an Amazon S3 bucket. The Security team now mandates that at-rest encryption be turned on immediately for all aspects of the application, without creating new resources and without any downtime.

To satisfy the requirements, which one of these services can the SysOps Administrator enable at-rest encryption on?

- A. EBS General Purpose SSD volumes
- B. RDS PostgreSQL database
- C. Amazon EFS file systems
- D. S3 objects within a bucket

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 549

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator noticed that a large number of Elastic IP addresses are being created on the company's AWS account, but they are not being associated with Amazon EC2 instances, and are incurring Elastic IP address charges in the monthly bill.

How can the Administrator identify who is creating the Elastic IP address?

- A. Attach a cost-allocation tag to each requested Elastic IP address with the IAM user name of the Developer who creates it.
- B. Query AWS CloudTrail logs by using Amazon Athena to search for Elastic IP address events.
- C. Create a CloudWatch alarm on the EIPCreated metric and send an Amazon SNS notification when the alarm triggers.
- D. Use Amazon Inspector to get a report of all Elastic IP addresses created in the last 30 days.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 550

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An application is running on Amazon EC2 instances behind a Classic Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. Occasionally multiple incoming requests will receive a 5xx HTTP response when making a request to the Classic Load Balancer. From the Amazon CloudWatch metrics, a SysOps Administrator observes the Elastic Load Balancing (ELB) SpillOverCount metric to be greater than zero during these occasions. These errors can be avoided by triggering scaling actions on which ELB metric?

- A. HealthyHostCount
- B. BackendConnectionErrors
- C. SurgeQueueLength
- D. UnHealthyHostCount

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 551

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An application running by a SysOps Administrator is under repeated, large-scale distributed denial of service (DDoS) attacks. Each time an attack occurs, multiple customers reach out to the Support team to report outages. The Administrator wants to minimize potential downtime from the DDoS attacks. The company requires 24/7 support.

Which AWS service should be set up to protect the application?

- A. AWS Trusted Advisor
- B. AWS Shield Advanced
- C. Amazon Cognito
- D. Amazon Inspector

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 552

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Malicious traffic is reaching company web servers from a single IP address located in another country. The SysOps Administrator is tasked with blocking this IP address.

How should the Administrator implement the restriction?

- A. Edit the security group for the web servers and add a deny entry for the IP address
- B. Edit the network access control list for the web server subnet and add a deny entry for the IP address
- C. Edit the VPC route table to route the malicious IP address to a black hole
- D. Use Amazon CloudFront's geo restriction feature to block traffic from the IP address

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 553

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator needs Amazon EC2 instances in two different VPCs in private subnets to be able to communicate. A peering connection between the two VPCs has been created using the AWS Management Console and shows a status of Active. The instances are still unable to send traffic to each other.

Why are the EC2 instances unable to communicate?

- A. One or both of the VPCs do not have an Internet Gateway attached
- B. The route tables have not been updated
- C. The peering connection has not been properly tagged
- D. One or both of the instances do not have an Elastic IP address assigned

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 554

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator must ensure that AWS CloudFormation deployment changes are properly tracked for governance.

Which AWS service should be used to accomplish this?

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

[Show Suggested Answer](#)





Exam question from Amazon's AWS-SysOps

Question #: 555

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

With the threat of ransomware viruses encrypting and holding company data hostage, which action should be taken to protect an Amazon S3 bucket?

- A. Deny Post, Put, and Delete on the bucket
- B. Enable server-side encryption on the bucket
- C. Enable Amazon S3 versioning on the bucket
- D. Enable snapshots on the bucket

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 556

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator has implemented an Auto Scaling group with a step scaling policy. The Administrator notices that the additional instances have not been included in the aggregated metrics. Why are the additional instances missing from the aggregated metrics?

- A. The warm-up period has not expired
- B. The instances are still in the boot process
- C. The instances have not been attached to the Auto Scaling group
- D. The instances are included in a different set of metrics

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 557

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Recently several critical files were mistakenly deleted from a shared Amazon S3 bucket. A SysOps Administrator needs to prevent accidental deletions from occurring in the future by enabling MFA Delete.

Once enabled, which bucket activities will require MFA authentication? (Choose two.)

- A. Permanently removing an object version from the bucket
- B. Disabling default object encryption for the bucket
- C. Listing all versions of deleted objects in the bucket
- D. Suspending versioning on the bucket
- E. Enabling MFA Add on the bucket

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 558

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator has an AWS Lambda function that stops all Amazon EC2 instances in a test environment at night and on the weekend. Stopping instances causes some servers to become corrupt due to the nature of the applications running on them.

What can the SysOps Administrator use to identify these EC2 instances?

- A. AWS Config
- B. Amazon EC2 termination protection
- C. Resource tagging
- D. Amazon CloudWatch

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 559

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company has Amazon EC2 instances that serve web content behind an Elastic Load Balancing (ELB) load balancer. The ELB Amazon CloudWatch metrics from a few hours ago indicate a significant number of 4XX errors. The EC2 instances from the time of these errors have been deleted.

At the time of the 4XX errors, how can an Administrator obtain information about who originated these requests?

- A. If ELB access logs have been enabled, the information can be retrieved from the S3 bucket
- B. Contact AWS Support to obtain application logs from the deleted instances
- C. Amazon S3 always keeps a backup of application logs from EC2 instances. Retrieve these logs for analysis
- D. Use AWS Trusted Advisor to obtain ELB access logs

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 560

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator is managing an application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS MySQL DB instance. The Administrator must ensure that that application stays available if the database becomes unresponsive.

How can these requirements be met?

- A. Create read replicas for the RDS database and use them in case of a database failure
- B. Create a new RDS instance from the snapshot of the original RDS instance if a failure occurs
- C. Keep a separate RDS database running and switch the endpoint in the web application if a failure occurs
- D. Modify the RDS instance to be a Multi-AZ deployment

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 561

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company has an asynchronous nightly process that feeds the results to a data warehouse system for weekly and monthly reporting. The process is running on a fleet of Amazon EC2 instances. A SysOps Administrator has been asked to identify ways to reduce the cost of running this process.

What is the MOST cost-effective solution?

- A. Use On-Demand EC2 instances in an Auto Scaling group
- B. Use Spot Instances to bid for the EC2 instances
- C. Use Reserved Instances to ensure the capacity
- D. Put the EC2 instances in a placement group

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 562

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A new website will run on Amazon EC2 instances behind an Application Load Balancer. Amazon Route 53 will be used to manage DNS records.

What type of record should be set in Route 53 to point the website's apex domain name (for example, `company.com`) to the Application Load Balancer?

- A. CNAME
- B. SOA
- C. TXT
- D. ALIAS

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 563

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An application running on Amazon EC2 allows users to launch batch jobs for data analysis. The jobs are run asynchronously, and the user is notified when they are complete. While multiple jobs can run concurrently, a user's request need not be fulfilled for up to 24 hours. To run a job, the application launches an additional EC2 instance that performs all the analytics calculations. A job takes between 75 and 110 minutes to complete and cannot be interrupted.

What is the MOST cost-effective way to run this workload?

- A. Run the application on On-Demand EC2 instances. Run the jobs on Spot Instances with a specified duration.
- B. Run the application on Reserved Instance EC2 instances. Run the jobs on AWS Lambda.
- C. Run the application on On-Demand EC2 instances. Run the jobs on On-Demand EC2 instances.
- D. Run the application on Reserved Instance EC2 instances. Run the jobs on Spot Instances with a specified duration.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 564

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A developer deploys an application running on Amazon EC2 by using an AWS CloudFormation template. The developer launches the stack from the console logged in as an AWS Identity and Access Management (IAM) user. When a SysOps Administrator attempts to run the same AWS CloudFormation template in the same AWS account from the console, it fails and returns the error:

```
`The image id `[ami-2a69aa47]' does not exist`
```

What is the MOST likely cause of the failure?

- A. The Administrator does not have the same IAM permissions as the developer.
- B. The Administrator used a different SSH key from that of the developer.
- C. The Administrator is running the template in a different region.
- D. The Administrator's Amazon EC2 service limits have been exceeded

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 565

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company has configured a library of IAM roles that grant access to various AWS resources. Each employee has an AWS IAM user, some of which have the permission to launch Amazon EC2 instances. The SysOps Administrator has attached the following policy to those users:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": ["ec2:*"],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": "iam:PassRole",
    "Resource": "arn:aws:iam::123456789012:role/InfraTeam*"
  }
]
```

What would be the result of this policy?

- A. Users are able to switch only to a role name that begins with `InfraTeam` followed by any other combination of characters.
- B. Users with the role of `InfraTeamLinux` are able to launch an EC2 instance and attach that role to it.
- C. `InfraTeam` role is being passed to a user who has full EC2 access.
- D. EC2 instances that are launched by these users have full AWS permissions.

Show Suggested Answer

Exam question from Amazon's AWS-SysOps

Question #: 566

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Application developers are reporting Access Denied errors when trying to list the contents of an Amazon S3 bucket by using the IAM user `arn:aws:iam::111111111111:user/application`. The following S3 bucket policy is in use:

```
{
  "Id": "S3BucketPolicy",
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "List",
      "Action": [
        "s3:List*"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:s3:::bucketname/*"
      ],
      "Principal": {
        "AWS": [
          "arn:aws:iam::111111111111:user/application"
        ]
      }
    }
  ]
}
```

How should a SysOps Administrator modify the S3 bucket policy to fix the issue?

- A. Change the `Effect` from `Allow` to `Deny`
- B. Change the `Action` from `s3:List*` to `s3:ListBucket`
- C. Change the `Resource` from `arn:aws:s3:::bucketname/*` to `arn:aws:s3:::bucketname`
- D. Change the `Principal` from `arn:aws:iam::111111111111:user/application` to `arn:aws:iam::111111111111:role/application`

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 567

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization has hired an external firm to audit unauthorized changes on the company's AWS environment, the external auditor needs appropriate access. How can this be accomplished?

- A. Create an IAM user and assign them a new policy with GetResources access on AWS Artifact
- B. Create an IAM user and add them to the existing "Administrator" IAM group
- C. Create an IAM user and assign them a new IAM policy with read access to the AWS CloudTrail logs in Amazon S3
- D. Create an IAM user and assign them a new policy with ListFindings access on Amazon Inspector

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 568

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator wants to automate the process of configuration, deployment, and management of Amazon EC2 instances using Chef or Puppet. Which AWS service will satisfy the requirement?

- A. AWS Elastic Beanstalk
- B. AWS CloudFormation
- C. AWS OpsWorks
- D. AWS Config

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 569

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A photo-sharing site delivers content worldwide from a library on Amazon S3 using Amazon CloudFront. Users are trying to access photos that either do not exist or they are not authorized to view.

What should be monitored to better understand the extent of this issue?

- A. GetRequests S3 metric on Amazon CloudWatch
- B. 4XXErrorRate CloudFront metric on CloudWatch
- C. 5XXErrorRate CloudFront metric on CloudWatch
- D. PostRequests S3 metric on CloudWatch

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 570

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company must share monthly report files that are uploaded to Amazon S3 with a third party. The third-party user list is dynamic, is distributed, and changes frequently. The least amount of access must be granted to the third party. Administrative overhead must be low for the internal teams who manage the process. How can this be accomplished while providing the LEAST amount of access to the third party?

- A. Allow only specified IP addresses to access the S3 buckets which will host files that need to be provided to the third party.
- B. Create an IAM role with the appropriate access to the S3 bucket, and grant login permissions to the console for the third party to access the S3 bucket.
- C. Create a pre-signed URL that can be distributed by email to the third party, allowing it to download specific S3 files.
- D. Have the third party sign up for an AWS account, and grant it cross-account access to the appropriate S3 bucket in the source account.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 571

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An administrator is responding to an alarm that reports increased application latency. Upon review, the Administrator notices that the Amazon RDS Aurora database frequently runs at 100% CPU utilization. The application is read heavy and does frequent lookups of a product table.

What should the Administrator do to reduce the application latency?

- A. Move the product table to Amazon Redshift and use an interleaved sort key
- B. Add Aurora Replicas and use a Reader Endpoint for product table lookups
- C. Move the product table to Amazon CloudFront and set the cache-control headers to public
- D. Use Auto Scaling to add extra Aurora nodes and set a trigger based on CPU utilization

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 572

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company is running a new promotion that will result in a massive spike in traffic for a single application. The SysOps Administrator must prepare the application and ensure that the customers have a great experience. The application is heavy on memory and is running behind an AWS Application Load Balancer (ALB).

The ALB has been pre-warmed, and the application is in an Auto Scaling group.

What built-in metric should be used to control the Auto Scaling group's scaling policy?

- A. RejectedConnection Count
- B. Request CountPerTarget
- C. CPUUtilization
- D. MemoryUtilization

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 573

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator is reviewing AWS Trusted Advisor warnings and encounters a warning for an S3 bucket policy that has open access permissions. While discussing the issue the bucket owner, the Administrator realizes the S3 bucket is an origin for an Amazon CloudFront web distribution.

Which action should the Administrator take to ensure that users access objects in Amazon S3 by using only CloudFront URLs?

- A. Encrypt the S3 bucket content with Server-Side Encryption with Amazon S3-Managed Keys (SSE-S3)
- B. Create an origin access identity and grant it permissions to read objects in the S3 bucket
- C. Assign an IAM user to the CloudFront distribution and whitelist the IAM user in the S3 bucket policy
- D. Assign an IAM role to the CloudFront distribution and whitelist the IAM role in the S3 bucket policy

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 574

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An e-commerce company hosts its website on the AWS us-west-1 region. It plans to create a special site for a promotion that should be visible only to shoppers from Canada.

What change should the SysOps Administrator make to the company's existing AWS setup to achieve this result?

- A. Update the Amazon Route 53 record set to use a latency routing policy for the new site
- B. Update the Application Load Balancer with a new host-based routing rule for the new site
- C. Update the Amazon Route 53 record set to use a geolocation routing policy for the new site
- D. Update the Application Load Balancer with a new path-based routing rule for the new site

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 575

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company currently has a single AWS account used by all project teams. The company is migrating to a multi-account strategy, where each project team will have its own account. The AWS IAM configuration must have the same roles and policies for each of the accounts.

What is the MOST efficient way to implement and manage these new requirements?

- A. Create a portfolio in the AWS Service Catalog for the IAM roles and policies. Have a specific product in the portfolio for each environment, project, and team that can be launched independently by each user.
- B. Use AWS Organizations to create organizational units (OUs) for each group of projects and each team. Then leverage service control policies at the account level to restrict what services can be used and what actions the users, groups, and roles can perform in those accounts.
- C. Create an AWS Lambda script that leverages cross-account access to each AWS account, and create all the roles and policies needed using the IAM API and JSON documents stored in Amazon S3.
- D. Create a single AWS CloudFormation template. Use CloudFormation StackSets to launch the CloudFormation template into each target account from the Administrator account.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 576

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator is creating an Amazon EC2 instance and has received an InsufficientInstanceCapacity error.

What is the cause of the error and how can it be corrected?

- A. AWS does not currently have enough capacity to service the request for that instance type. A different Availability Zone or instance type must be used.
- B. The account has reached its concurrent running instance limit. An EC2 limit increase request must be filed with AWS Support.
- C. The APIs that service the EC2 requests have received too many requests and capacity has been reached. The request should be attempted again in a few minutes.
- D. The Administrator did not specify the correct size of the instance to support the capacity requirements of the workload. Select a bigger instance.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 577

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A web application runs on Amazon EC2 instances with public IPs assigned behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS Multi-AZ DB instance. The Application Load Balancer, EC2 instances, and RDS DB instance all run in separate sets of subnets. The EC2 instances can communicate with the DB instance, but cannot connect with external services.

What is the MOST likely solution?

- A. Assign a public IP address to the database server and restart the database engine.
- B. Create and attach an Internet gateway to the VPC. Create a route table for the EC2 instance's subnets that sends Internet traffic to the gateway.
- C. Create and attach a virtual private gateway to the VPC. Create a route table for the EC2 instances' subnets that sends Internet traffic to the gateway.
- D. Create a VPC peering connection to a VPC that has an Internet gateway attached. Create a route table for the EC2 instances' subnets that sends Internet traffic to the peered VPC.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 578

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company has deployed a new application running on Amazon EC2 instances. The application team must verify for the Security team that all common vulnerabilities and exposures have been addressed, both now and regularly throughout the application's lifespan.

How can the Application team satisfy the Security team's requirement?

- A. Perform regular assessments with Amazon Inspector
- B. Perform regular assessments with AWS Trusted Advisor
- C. Integrate AWS Personal Health Dashboard with Amazon CloudWatch events to get security notifications
- D. Grant the Administrator and Security team access to AWS Artifact

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 579

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

InfoSec is concerned that an employee may expose sensitive data in an Amazon S3 bucket.

How can this concern be addressed without putting undue restrictions on users?

- A. Apply an IAM policy on all users that denies the action s3:PutBucketPolicy
- B. Restrict S3 bucket access to specific IAM roles managed using federated access
- C. Activate an AWS Config rule to identify public buckets and alert InfoSec using Amazon SNS
- D. Email the findings of AWS Personal Health Dashboard to InfoSec daily

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 580

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator is using AWS CloudFormation to deploy resources but would like to manually address any issues that the template encounters. What should the Administrator add to the template to support the requirement?

- A. Enable Termination Protection on the stack
- B. Set the OnFailure parameter to `DO_NOTHING`
- C. Restrict the IAM permissions for CloudFormation to delete resources
- D. Set the DeleteStack API action to `No`

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 581

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator discovers the organization's tape archival system is no longer functioning in its on-premises data center.

What AWS service can be used to create a virtual tape interface to replace the physical tape system?

- A. AWS Snowball
- B. AWS SMS
- C. Amazon Glacier
- D. AWS Storage Gateway

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 582

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A new application runs on Amazon EC2 instances and accesses data in an Amazon RDS database instance. When fully deployed in production, the application fails. The database can be queried from a console on a bastion host. When looking at the web server logs, the following error is repeated multiple times:

*** Error Establishing a Database Connection.

Which of the following may be causes of the connectivity problems? (Choose two.)

- A. The security group for the database does not have the appropriate egress rule from the database to the web server.
- B. The certificate used by the web server is not trusted by the RDS instance.
- C. The security group for the database does not have the appropriate ingress rule from the web server to the database.
- D. The database is still being created and is not available for connectivity.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 583

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A recent audit found that most resources belonging to the Development team were in violation of patch compliance standards. The resources were properly tagged. Which service should be used to quickly remediate the issue and bring the resources back into compliance?

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 584

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator has been able to consolidate multiple, secure websites onto a single server, and each site is running on a different port. The Administrator now wants to start a duplicate server in a second Availability Zone and put both behind a load balancer for high availability.

What would be the command line necessary to deploy one of the sites' certificates to the load balancer?

A.

```
aws kms modify-listener --load-balancer-name my-load-balancer - -certificates CertificateArn=arn:aws:iam::123456789012:server-certificate/my-new-server-cert
```

B.

```
aws elb set-load-balancer-listener-ssl-certificate - -load-balancer-name my-load-balancer - -load-balancer-port 443 - -ssl-certificate-id arn:aws:iam::123456789012:server-certificate/new-server-cert
```

C.

```
aws ec2 put-ssl-certificate - -load-balancer-name my-load-balancer - -load-balancer-port 443 - -ssl-certificate-id arn:aws:iam::123456789012:server-certificate/new-server-cert
```

D.

```
aws acm put-ssl-certificate - -load-balancer-name my-load-balancer- -load-balancer-port 443 - -ssl-certificate-id arn:aws:iam::123456789012:server-certificate/new-server-cert
```

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 585

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An Amazon EBS volume attached to an EC2 instance was recently modified. Part of the modification included increasing the storage capacity. The SysOps Administrator notices that the increased storage capacity is not reflected in the file system.

Which step should the Administrator complete to use the increased storage capacity?

- A. Restart the EC2 instance.
- B. Extend the volume's file system.
- C. Detach the EBS volume, resize it, and attach it.
- D. Take an EBS snapshot and restore it to the bigger volume.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 586

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator is creating additional Amazon EC2 instances and receives an InstanceLimitExceeded error.

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

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 587

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator is troubleshooting Amazon EC2 connectivity issues to the internet. The EC2 instance is in a private subnet. Below is the route table that is applied to the subnet of the EC2 instance.

Destination ` 10.2.0.0/16 -

Target ` local -

Status ` Active -

Propagated ` No -

Destination ` 0.0.0.0/0 -

Target ` nat-xxxxxxx -

Status ` Blackhole -

Propagated ` No -

What has caused the connectivity issue?

- A. The NAT gateway no longer exists.
- B. There is no route to the internet gateway.
- C. The routes are no longer propagating.
- D. There is no route rule with a destination for the internet.

Show Suggested Answer

Exam question from Amazon's AWS-SysOps

Question #: 588

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

Malicious traffic is reaching company web servers. A SysOps Administrator is tasked with blocking this traffic. The malicious traffic is distributed over many IP addresses and represents much higher traffic than is typically seen from legitimate users.

How should the Administrator protect the web servers?

- A. Create a security group for the web servers and add deny rules for malicious sources.
- B. Set the network access control list for the web servers' subnet and add deny entries.
- C. Place web servers behind AWS WAF and establish the rate limit to create a blacklist.
- D. Use Amazon CloudFront to cache all pages and remove the traffic from the web servers.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 589

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator must evaluate storage solutions to replace a company's current user-shared drives infrastructure. Any solution must support security controls that enable Portable Operating System Interface (POSIX) permissions and Network File System protocols. Additionally, any solution must be accessible from multiple Amazon EC2 instances and on-premises servers connected to the Amazon VPC.

Which AWS service meets the user drive requirements?

- A. Amazon S3
- B. Amazon EFS
- C. Amazon EBS
- D. Amazon SQS

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 590

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A Developer created an AWS Lambda function and has asked the SysOps Administrator to make this function run every 15 minutes. What is the MOST efficient way to accomplish this request?

- A. Create an Amazon EC2 instance and schedule a cron to invoke the Lambda function.
- B. Create a Repeat Time variable inside the Lambda function to invoke the Lambda function.
- C. Create a second Lambda function to monitor and invoke the first Lambda function.
- D. Create an Amazon CloudWatch scheduled event to invoke the Lambda function.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 591

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company's auditor implemented a compliance requirement that all Amazon S3 buckets must have logging enabled. A SysOps administrator is tasked to ensure this compliance requirement is met, while still permitting developers to create and use new S3 buckets.

Which action should be taken to accomplish this?

- A. Add AWS CloudTrail logging for the S3 buckets.
- B. Implement IAM policies to allow only the storage team to create S3 buckets.
- C. Add the S3_BUCKET_LOGGING_ENABLED AWS Config managed rule.
- D. Create an AWS Lambda function to delete the S3 buckets if logging is not turned on.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 592

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization is concerned that its Amazon RDS databases are not protected. The solution to address this issue must be low cost, protect against table corruption that could be overlooked for several days, and must offer a 30-day window of protection.

How can these requirements be met?

- A. Enable Multi-AZ on the RDS instance to maintain the data in a second Availability Zone.
- B. Create a read replica of the RDS instance to maintain the data in a second region.
- C. Ensure that automated backups are enabled and set the appropriate retention period.
- D. Enable versioning in RDS to recover altered table data when needed.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 593

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization is running multiple applications for their customers. Each application is deployed by running a base AWS CloudFormation template that configures a new VPC. All applications are run in the same AWS account and AWS Region. A SysOps Administrator has noticed that when trying to deploy the same AWS CloudFormation stack, it fails to deploy.

What is likely to be the problem?

- A. The Amazon Machine image used is not available in that region.
- B. The AWS CloudFormation template needs to be updated to the latest version.
- C. The VPC configuration parameters have changed and must be updated in the template.
- D. The account has reached the default limit for VPCs allowed.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 594

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

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

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 595

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A web-commerce application stores its data in an Amazon Aurora DB cluster with an Aurora replica. The application displays shopping cart information by reading data from the reader endpoint. When monitoring the Aurora database, the SysOps Administrator sees that the AuroraReplicaLagMaximum metric for a single replica is high. What behavior is the application MOST likely exhibiting to users?

- A. Users cannot add any items to the shopping cart.
- B. Users intermittently notice that the cart is not updated correctly.
- C. Users cannot remove any items from the shopping cart.
- D. Users cannot use the application because it is falling back to an error page.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 596

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company would like to review each change in the infrastructure before deploying updates in its AWS CloudFormation stacks.

Which action will allow an Administrator to understand the impact of these changes before implementation?

- A. Implement a blue/green strategy using AWS Elastic Beanstalk.
- B. Perform a canary deployment using Application Load Balancers and target groups.
- C. Create a change set for the running stack.
- D. Submit the update using the UpdateStack API call.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 597

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A Systems Administrator is responsible for maintaining custom, approved AMIs for a company. These AMIs must be shared with each of the company's AWS accounts.

How can the Administrator address this issue?

- A. Contact AWS Support for sharing AMIs with other AWS accounts.
- B. Modify the permissions on the AMIs so that they are publicly accessible.
- C. Modify the permissions on the IAM role that are associated with the AMI.
- D. Share the AMIs with each AWS account using the console or CLI.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 598

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company's data retention policy dictates that backups be stored for exactly two years. After that time, the data must be deleted.

How can Amazon EBS snapshots be managed to conform to this data retention policy?

- A. Use an Amazon S3 lifecycle policy to delete snapshots older than two years.
- B. Configure Amazon Inspector to find and delete old EBS snapshots.
- C. Schedule an AWS Lambda function using Amazon CloudWatch Events to periodically run a script to delete old snapshots.
- D. Configure an Amazon CloudWatch alarm to trigger the launch of an AWS CloudFormation template that will clean the older snapshots.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 599

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator must devise a strategy for enforcing tagging of all EC2 instances and Amazon Elastic Block Store (Amazon EBS) volumes. What action can the Administrator take to implement this for real-time enforcement?

- A. Use the AWS Tag Editor to manually search for untagged resources and then tag them properly in the editor.
- B. Set up AWS Service Catalog with the TagOptions Library rule that enforces a tagging taxonomy proactively when instances and volumes are launched.
- C. In a PowerShell or shell script, check for untagged items by using the resource tagging GetResources API action, and then manually tag the reported items.
- D. Launch items by using the AWS API. Use the TagResources API action to apply the required tags when the instances and volumes are launched.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 600

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

During a security investigation, it is determined that there is a coordinated attack on the web applications deployed on Amazon EC2. The attack is performed through malformed HTTP headers.

What AWS service or feature would prevent this traffic from reaching the EC2 instances?

- A. Amazon Inspector
- B. Amazon Security Groups
- C. AWS WAF
- D. Application Load Balancer (ALB)

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 601

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company is deploying a legacy web application on Amazon EC2 instances behind an ELB Application Load Balancer. The application worked well in the test environment. However, in production, users report that they are prompted to log in to the system several times an hour. Which troubleshooting step should be taken to help resolve the problem reported by users?

- A. Confirm that the Application Load Balancer is in a multi-AZ configuration.
- B. Enable health checks on the Application Load Balancer.
- C. Ensure that port 80 is configured on the security group.
- D. Enable sticky sessions on the Application Load Balancer.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 602

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company has mandated the use of multi-factor authentication (MFA) for all IAM users, and requires users to make all API-calls using the CLI. However, users are not prompted to enter MFA tokens, and are able to run CLI commands without MFA. In an attempt to enforce MFA, the company attached an IAM policy to all users that denies API calls that have not been authenticated with MFA.

What additional step must be taken to ensure that API calls are authenticated using MFA?

- A. Enable MFA on IAM roles, and require IAM users to use role credentials to sign API calls.
- B. Ask the IAM users to log into the AWS Management Console with MFA before making API calls using the CLI.
- C. Restrict the IAM users to use of the console, as MFA is not supported for CLI use.
- D. Require users to use temporary credentials from the get-session token command to sign API calls.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 603

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An application is being developed that will be served across a fleet of Amazon EC2 instances, which require a consistent view of persistent data. Items stored vary in size from 1KB to 300MB; the items are read frequently, created occasionally, and often require partial changes without conflict. The data store is not expected to grow beyond 2TB, and items will be expired according to age and content type.

Which AWS service solution meets these requirements?

- A. Amazon S3 buckets with lifecycle policies to delete old objects.
- B. Amazon RDS PostgreSQL and a job that deletes rows based on age and file type columns.
- C. Amazon EFS and a scheduled process to delete files based on age and extension.
- D. An EC2 instance store synced on boot from a central Amazon EBS-backed instance.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 604

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator created an Amazon VPC with an IPv6 CIDR block, which requires access to the internet. However, access from the internet towards the VPC is prohibited. After adding and configuring the required components to the VPC, the Administrator is unable to connect to any of the domains that reside on the internet.

What additional route destination rule should the Administrator add to the route tables?

- A. Route `::/0` traffic to a NAT gateway
- B. Route `::/0` traffic to an internet gateway
- C. Route `0.0.0.0/0` traffic to an egress-only internet gateway
- D. Route `::/0` traffic to an egress-only internet gateway

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 605

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A recent organizational audit uncovered an existing Amazon RDS database that is not currently configured for high availability. Given the critical nature of this database, it must be configured for high availability as soon as possible.

How can this requirement be met?

- A. Switch to an active/passive database pair using the `create-db-instance-read-replica` with the `--availability-zone` flag.
- B. Specify high availability when creating a new RDS instance, and live-migrate the data.
- C. Modify the RDS instance using the console to include the Multi-AZ option.
- D. Use the `modify-db-instance` command with the `--ha` flag.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 606

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company must ensure that any objects uploaded to an S3 bucket are encrypted.

Which of the following actions will meet this requirement? (Choose two.)

- A. Implement AWS Shield to protect against unencrypted objects stored in S3 buckets.
- B. Implement Object access control list (ACL) to deny unencrypted objects from being uploaded to the S3 bucket.
- C. Implement Amazon S3 default encryption to make sure that any object being uploaded is encrypted before it is stored.
- D. Implement Amazon Inspector to inspect objects uploaded to the S3 bucket to make sure that they are encrypted.
- E. Implement S3 bucket policies to deny unencrypted objects from being uploaded to the buckets.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 607

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

When the AWS Cloud infrastructure experiences an event that may impact an organization, which AWS service can be used to see which of the organization's resources are affected?

- A. AWS Service Health Dashboard
- B. AWS Trusted Advisor
- C. AWS Personal Health Dashboard
- D. AWS Systems Manager

[Show Suggested Answer](#)



Exam question from Amazon's AWS-SysOps

Question #: 608

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company's static website hosted on Amazon S3 was launched recently, and is being used by tens of thousands of users. Subsequently, website users are experiencing 503 service unavailable errors.

Why are these errors occurring?

- A. The request rate to Amazon S3 is too high.
- B. There is an error with the Amazon RDS database.
- C. The requests to Amazon S3 do not have the proper permissions.
- D. The users are in a different geographical region and Amazon Route 53 is restricting access.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 609

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An organization has two AWS accounts: Development and Production. A SysOps Administrator manages access of IAM users to both accounts. Some IAM users in Development should have access to certain resources in Production.

How can this be accomplished?

- A. Create an IAM role in the Production account with the Development account as a trusted entity and then allow those users from the Development account to assume the Production account IAM role.
- B. Create a group of IAM users in the Development account, and add Production account service ARNs as resources in the IAM policy.
- C. Establish a federation between the two accounts using the on-premises Microsoft Active Directory, and allow the Development account to access the Production account through this federation.
- D. Establish an Amazon Cognito Federated Identity between the two accounts, and allow the Development account to access the Production account through this federation.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 610

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator is responsible for managing a set of 12.micro Amazon EC2 instances. The Administrator wants to automatically reboot any instance that exceeds 80% CPU utilization.

Which of these solutions would meet the requirements?

- A. Create an Amazon CloudWatch alarm on the CPUCreditBalance metric and specify a terminate alarm action.
- B. Create an Amazon CloudWatch alarm on the CPUUtilization metric and specify a reboot alarm action.
- C. Create an Amazon CloudWatch alarm on the CPUCreditBalance metric and specify a reboot alarm action.
- D. Create an Amazon CloudWatch alarm on the CPUUtilization metric and specify a terminate alarm action.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 611

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company's customers are reporting increased latency while accessing static web content from Amazon S3. A SysOps Administrator observed a very high rate of read operations on a particular S3 bucket.

What will minimize latency by reducing load on the S3 bucket?

- A. Migrate the S3 bucket to a region that is closer to end users' geographic locations.
- B. Use cross-region replication to replicate all of the data to another region.
- C. Create an Amazon CloudFront distribution with the S3 bucket as the origin.
- D. Use Amazon ElastiCache to cache data being served from Amazon S3.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 612

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company requires that all access from on-premises applications to AWS services go over its AWS Direct Connect connection rather than the public internet. How would a SysOps Administrator implement this requirement?

- A. Implement an IAM policy that uses the `aws:sourceConnection` condition to allow access from the AWS Direct Connect connection ID only
- B. Set up a public virtual interface on the AWS Direct Connect connection
- C. Configure AWS Shield to protect the AWS Management Console from being accessed by IP addresses other than those within the data center ranges
- D. Update all the VPC network ACLs to allow access from the data center IP ranges

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 613

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator must find a way to set up alerts when Amazon EC2 service limits are close to being reached.

How can the Administrator achieve this requirement?

- A. Use Amazon Inspector and Amazon CloudWatch Events.
- B. Use AWS Trusted Advisor and Amazon CloudWatch Events.
- C. Use the Personal Health Dashboard and CloudWatch Events.
- D. Use AWS CloudTrail and CloudWatch Events.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 614

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A web application accepts orders from online users and places the orders into an Amazon SQS queue. Amazon EC2 instances in an EC2 Auto Scaling group read the messages from the queue, process the orders, and email order confirmations to the users. The Auto Scaling group scales up and down based on the queue depth. At the beginning of each business day, users report confirmation emails are delayed.

What action will address this issue?

- A. Create a scheduled scaling action to scale up in anticipation of the traffic.
- B. Change the Auto Scaling group to scale up and down based on CPU utilization.
- C. Change the launch configuration to launch larger EC2 instance types.
- D. Modify the scaling policy to deploy more EC2 instances when scaling up.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 615

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company creates custom AMI images by launching new Amazon EC2 instances from an AWS CloudFormation template. It installs and configures necessary software through AWS OpsWorks, and takes images of each EC2 instance. The process of installing and configuring software can take between 2 to 3 hours, but at times, the process stalls due to installation errors.

The SysOps Administrator must modify the CloudFormation template so if the process stalls, the entire stack will fail and roll back.

Based on these requirements, what should be added to the template?

- A. Conditions with a timeout set to 4 hours.
- B. CreationPolicy with a timeout set to 4 hours.
- C. DependsOn with a timeout set to 4 hours.
- D. Metadata with a timeout set to 4 hours.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 616

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator must take a team's single existing AWS CloudFormation template and split it into smaller, service-specific templates. All of the services in the template reference a single, shared Amazon S3 bucket.

What should the Administrator do to ensure that this S3 bucket can be referenced by all the service templates?

- A. Include the S3 bucket as a mapping in each template.
- B. Add the S3 bucket as a resource in each template.
- C. Create the S3 bucket in its own template and export it.
- D. Generate the S3 bucket using StackSets.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 617

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

After installing and configuring the Amazon CloudWatch agent on an EC2 instance, the anticipated system logs are not being received by CloudWatch Logs. Which of the following are likely to be the cause of this problem? (Choose two.)

- A. A custom of third-party solution for logs is being used.
- B. The IAM role attached to the EC2 instance does not have the proper permissions.
- C. The CloudWatch agent does not support the operating system used.
- D. A billing constraint is limiting the number of CloudWatch Logs within this account.
- E. The EC2 instance is in a private subnet, and the VPC does not have a NAT gateway.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 618

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator found that a newly-deployed Amazon EC2 application server is unable to connect to an existing Amazon RDS database. After enabling VPC Flow Logs and confirming that the flow log is active on the console, the log group cannot be located in Amazon CloudWatch.

What are the MOST likely reasons for this situation? (Choose two.)

- A. The Administrator must configure the VPC Flow Logs to have them sent to AWS CloudTrail.
- B. The Administrator has waited less than ten minutes for the log group to be created in CloudWatch.
- C. The account VPC Flow Logs have been disabled by using a service control policy.
- D. No relevant traffic has been sent since the VPC Flow Logs were created
- E. The account has Amazon GuardDuty enabled.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 619

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An HTTP web application is launched on Amazon EC2 instances behind an ELB Application Load Balancer. The EC2 instances run across multiple Availability Zones. A network ACL and a security group for the load balancer and EC2 instances allow inbound traffic on port 80. After launch, the website cannot be reached over the internet.

What additional step should be taken?

- A. Add a rule to the security group allowing outbound traffic on port 80.
- B. Add a rule to the network ACL allowing outbound traffic on port 80.
- C. Add a rule to the security group allowing outbound traffic on ports 1024 through 65535.
- D. Add a rule to the network ACL allowing outbound traffic on ports 1024 through 65535.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 620

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A company has an application that is running on an EC2 instance in one Availability Zone. A SysOps Administrator has been tasked with making the application highly available. The Administrator created a launch configuration from the running EC2 instance. The Administrator also properly configured a load balancer.

What step should the Administrator complete next to make the application highly available?

- A. Create an Auto Scaling group by using the launch configuration across at least 2 Availability Zones with a minimum size of 1, desired capacity of 1, and a maximum size of 1.
- B. Create an Auto Scaling group by using the launch configuration across at least 3 Availability Zones with a minimum size of 2, desired capacity of 2, and a maximum of 2.
- C. Create an Auto Scaling group by using the launch configuration across at least 2 regions with a minimum size of 1, desired capacity of 1, and a maximum size of 1.
- D. Create an Auto Scaling group by using the launch configuration across at least 3 regions with a minimum size of 2, desired capacity of 2, and a maximum size of 2.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 621

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

An Applications team has successfully deployed an AWS CloudFormation stack consisting of 30 t2-medium Amazon EC2 instances in the us-west-2 Region. When using the same template to launch a stack in us-east-2, the launch failed and rolled back after launching only 10 EC2 instances. What is a possible cause of this failure?

- A. The IAM user did not have privileges to launch the CloudFormation template.
- B. The t2.medium EC2 instance service limit was reached.
- C. An AWS Budgets threshold was breached.
- D. The application's Amazon Machine Image (AMI) is not available in us-east-2.

Show Suggested Answer





Exam question from Amazon's AWS-SysOps

Question #: 622

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

A SysOps Administrator stores crash dump files in Amazon S3. New security and privacy measures require that crash dumps older than 6 months be deleted. Which approach meets this requirement?

- A. Use Amazon CloudWatch Events to delete objects older than 6 months.
- B. Implement lifecycle policies to delete objects older than 6 months.
- C. Use the Amazon S3 Standard-Infrequent Access (S3 Standard-IA) storage class to automatically delete objects older than 6 months.
- D. Create versioning rules to delete objects older than 6 months.

Show Suggested Answer



Exam question from Amazon's AWS-SysOps

Question #: 623

Topic #: 1

[\[All AWS-SysOps Questions\]](#)

The Accounting department would like to receive billing updates more than once a month. They would like the updates to be in a format that can easily be viewed with a spreadsheet application.

How can this request be fulfilled?

- A. Use Amazon CloudWatch Events to schedule a billing inquiry on a bi-weekly basis. Use AWS Glue to convert the output to CSV.
- B. Set AWS Cost and Usage Reports to publish bills daily to an Amazon S3 bucket in CSV format.
- C. Use the AWS CLI to output billing data as JSON. Use Amazon SES to email bills on a daily basis.
- D. Use AWS Lambda, triggered by CloudWatch, to query billing data and push to Amazon RDS.

Show Suggested Answer

