

# Amazon

## Exam Questions AWS-SysOps

Amazon AWS Certified SysOps Administrator - Associate



### NEW QUESTION 1

- (Exam Topic 2)

You need to update an existing AWS CloudFormation stack. If needed, a copy to the CloudFormation template is available in an Amazon S3 bucket named cloudformation-bucket

- \* 1. Use the us-east-2 Region for all resources.
- \* 2. Unless specified below, use the default configuration settings.
- \* 3. update the Amazon EC2 instance named Devinstance by making the following changes to the stack named 1700182:
  - \* a) Change the EC2 instance type to us-east-t2.nano.
  - \* b) Allow SSH to connect to the EC2 instance from the IP address range 192.168.100.0/30.
  - \* c) Replace the instance profile IAM role with lamRoleB.
- \* 4. Deploy the changes by updating the stack using the CFServiceR01e role.
- \* 5. Edit the stack options to prevent accidental deletion.
- \* 6. Using the output from the stack, enter the value of the ProdInstanceid in the text box below:

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Here are the steps to update an existing AWS CloudFormation stack:

- Log in to the AWS Management Console and navigate to the CloudFormation service in the us-east-2 Region.
- Find the existing stack named 1700182 and click on it.
- Click on the "Update" button.
- Choose "Replace current template" and upload the updated CloudFormation template from the Amazon S3 bucket named "cloudformation-bucket"
- In the "Parameter" section, update the EC2 instance type to us-east-t2.nano and add the IP address range 192.168.100.0/30 for SSH access.
- Replace the instance profile IAM role with lamRoleB.
- In the "Capabilities" section, check the checkbox for "IAM Resources"
- Choose the role CFServiceR01e and click on "Update Stack"
- Wait for the stack to be updated.
- Once the update is complete, navigate to the stack and click on the "Stack options" button, and select "Prevent updates to prevent accidental deletion"
- To get the value of the ProdInstanceid, navigate to the "Outputs" tab in the CloudFormation stack and find the key "ProdInstanceid". The value corresponding to it is the value that you need to enter in the text box below.

Note:

- You can use AWS CloudFormation to update an existing stack.

You can use the AWS CloudFormation service role to deploy updates.

You can refer to the AWS CloudFormation documentation for more information on how to update and manage stacks: <https://aws.amazon.com/cloudformation/>

### NEW QUESTION 2

- (Exam Topic 1)

A company needs to restrict access to an Amazon S3 bucket to Amazon EC2 instances in a VPC only. All traffic must be over the AWS private network.

What actions should the SysOps administrator take to meet these requirements?

- A. Create a VPC endpoint for the S3 bucket, and create an IAM policy that conditionally limits all S3 actions on the bucket to the VPC endpoint as the source.
- B. Create a VPC endpoint for the S3 bucket, and create an S3 bucket policy that conditionally limits all S3 actions on the bucket to the VPC endpoint as the source.
- C. Create a service-linked role for Amazon EC2 that allows the EC2 instances to interact directly with Amazon S3, and attach an IAM policy to the role that allows the EC2 instances full access to the S3 bucket.
- D. Create a NAT gateway in the VPC, and modify the VPC route table to route all traffic destined for Amazon S3 through the NAT gateway.

**Answer:** B

#### Explanation:

While IAM policy (letter A) also can be used, it does not enforce everyone. The only option that enforces everyone is policy configured directly in the bucket S3.

### NEW QUESTION 3

- (Exam Topic 1)

A SysOps administrator is deploying an application on 10 Amazon EC2 instances. The application must be highly available. The instances must be placed on distinct underlying hardware.

What should the SysOps administrator do to meet these requirements?

- A. Launch the instances into a cluster placement group in a single AWS Region.
- B. Launch the instances into a partition placement group in multiple AWS Regions.
- C. Launch the instances into a spread placement group in multiple AWS Regions.
- D. Launch the instances into a spread placement group in single AWS Region

**Answer:** D

#### Explanation:

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

#### NEW QUESTION 4

- (Exam Topic 1)

A company is deploying a third-party unit testing solution that is delivered as an Amazon EC2 Amazon Machine Image (AMI). All system configuration data is stored in Amazon DynamoDB. The testing results are stored in Amazon S3.

A minimum of three EC2 instances are required to operate the product. The company's testing team wants to use an additional three EC2 Instances when the Spot Instance prices are at a certain threshold. A SysOps administrator must Implement a highly available solution that provides this functionality.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Define an Amazon EC2 Auto Scaling group by using a launch configuratio
- B. Use the provided AMI In the launch configuratio
- C. Configure three On-Demand Instances and three Spot Instance
- D. Configure a maximum Spot Instance price In the launch configuration.
- E. Define an Amazon EC2 Auto Scaling group by using a launch templat
- F. Use the provided AMI in the launch templat
- G. Configure three On-Demand Instances and three Spot Instance
- H. Configure a maximum Spot Instance price In the launch template.
- I. Define two Amazon EC2 Auto Scaling groups by using launch configuration
- J. Use the provided AMI in the launch configuration
- K. Configure three On-Demand Instances for one Auto Scaling grou
- L. Configure three Spot Instances for the other Auto Scaling grou
- M. Configure a maximum Spot Instance price in the launch configuration for the Auto Scaling group that has Spot Instances.
- N. Define two Amazon EC2 Auto Scaling groups by using launch template
- O. Use the provided AMI in the launch template
- P. Configure three On-DemandInstances for one Auto Scaling grou
- Q. Configure three Spot Instances for the other Auto Scaling grou
- R. Configure a maximum Spot Instance price in the launch template for the Auto Scaling group that has Spot Instances.

**Answer:** A

#### Explanation:

- > <https://docs.aws.amazon.com/autoscaling/ec2/userguide/LaunchTemplates.html>
- > <https://docs.aws.amazon.com/autoscaling/ec2/userguide/LaunchConfiguration.html>

#### NEW QUESTION 5

- (Exam Topic 1)

A company hosts a website on multiple Amazon EC2 instances that run in an Auto Scaling group. Users are reporting slow responses during peak times between 6 PM and 11 PM every weekend. A SysOps administrator must implement a solution to improve performance during these peak times.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create a scheduled Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function to increase the desired capacity before peak times.
- B. Configure a scheduled scaling action with a recurrence option to change the desired capacity before and after peak times.
- C. Create a target tracking scaling policy to add more instances when memory utilization is above 70%.
- D. Configure the cooldown period for the Auto Scaling group to modify desired capacity before and after peak times.

**Answer:** B

#### Explanation:

"Scheduled scaling helps you to set up your own scaling schedule according to predictable load changes. For example, let's say that every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can configure a schedule for Amazon EC2 Auto Scaling to increase capacity on Wednesday and decrease capacity on Friday." [https://docs.aws.amazon.com/autoscaling/ec2/userguide/schedule\\_time.html](https://docs.aws.amazon.com/autoscaling/ec2/userguide/schedule_time.html)

#### NEW QUESTION 6

- (Exam Topic 1)

A company has an Auto Scaling group of Amazon EC2 instances that scale based on average CPU utilization. The Auto Scaling group events log indicates an InsufficientInstanceCapacity error.

Which actions should a SysOps administrator take to remediate this issue? (Select TWO).

- A. Change the instance type that the company is using.
- B. Configure the Auto Scaling group in different Availability Zones.
- C. Configure the Auto Scaling group to use different Amazon Elastic Block Store (Amazon EBS) volume sizes.
- D. Increase the maximum size of the Auto Scaling group.
- E. Request an increase in the instance service quota.

**Answer:** AB

#### NEW QUESTION 7

- (Exam Topic 1)

A SysOps administrator creates an AWS CloudFormation template to define an application stack that can be deployed in multiple AWS Regions.

The SysOps administrator also creates an Amazon CloudWatch dashboard by using the AWS Management Console. Each deployment of the application requires its own CloudWatch dashboard.

How can the SysOps administrator automate the creation of the CloudWatch dashboard each time the application is deployed?

- A. Create a script by using the AWS CLI to run the aws cloudformation put-dashboard command with the name of the dashboard
- B. Run the command each time a new CloudFormation stack is created.
- C. Export the existing CloudWatch dashboard as JSON
- D. Update the CloudFormation template to define an AWS::CloudWatch::Dashboard resource
- E. Include the exported JSON in the resource's DashboardBody property.
- F. Update the CloudFormation template to define an resource

- G. Use the intrinsic Ref function to reference the ID of the existing CloudWatch dashboard.
- H. Update the CloudFormation template to define an AWS::CloudWatch::Dashboard resource.
- I. Specify the name of the existing dashboard in the DashboardName property.

**Answer: B**

**Explanation:**

You can only use the Intrinsic Ref function to reference a resource that is being created at the same time as the current CloudFormation template. The question states that the CloudWatch dashboard was previously created using the AWS Management Console, so there is no ID to reference the existing CloudWatch dashboard in the CloudFormation template. You would need to export the existing CloudWatch dashboard as JSON, then use the DashboardBody property in the CloudFormation template to replicate it upon each deployment  
(<https://docs.aws.amazon.com/AmazonCloudWatch/latest/APIReference/CloudWatch-Dashboard-Body-Structu>)

**NEW QUESTION 8**

- (Exam Topic 1)

A company has a memory-intensive application that runs on a fleet of Amazon EC2 instances behind an Elastic Load Balancer (ELB). The instances run in an Auto Scaling group. A SysOps administrator must ensure that the application can scale based on the number of users that connect to the application. Which solution will meet these requirements?

- A. Create a scaling policy that will scale the application based on the ActiveConnectionCount Amazon CloudWatch metric that is generated from the ELB.
- B. Create a scaling policy that will scale the application based on the mem used Amazon CloudWatch metric that is generated from the ELB.
- C. Create a scheduled scaling policy to increase the number of EC2 instances in the Auto Scaling group to support additional connections.
- D. Create and deploy a script on the ELB to expose the number of connected users as a custom Amazon CloudWatch metric.
- E. Create a scaling policy that uses the metric.

**Answer: D**

**Explanation:**

This solution will allow the application to scale based on the number of users that connect to the application. The other solutions (creating a scaling policy that uses the ActiveConnectionCount Amazon CloudWatch metric generated from the ELB, creating a scaling policy that uses the mem used Amazon CloudWatch metric generated from the ELB, or creating a scheduled scaling policy to increase the number of EC2 instances in the Auto Scaling group to support additional connections) will not meet the requirements, as they do not allow the application to scale based on the number of users that connect to the application.

**NEW QUESTION 9**

- (Exam Topic 1)

A SysOps administrator developed a Python script that uses the AWS SDK to conduct several maintenance tasks. The script needs to run automatically every night. What is the MOST operationally efficient solution that meets this requirement?

- A. Convert the Python script to an AWS Lambda function.
- B. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the function every night.
- C. Convert the Python script to an AWS Lambda function.
- D. Use AWS CloudTrail to invoke the function every night.
- E. Deploy the Python script to an Amazon EC2 instance.
- F. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the instance to start and stop every night.
- G. Deploy the Python script to an Amazon EC2 instance.
- H. Use AWS Systems Manager to schedule the instance to start and stop every night.

**Answer: A**

**NEW QUESTION 10**

- (Exam Topic 1)

A SysOps administrator is tasked with deploying a company's infrastructure as code. The SysOps administrator wants to write a single template that can be reused for multiple environments. How should the SysOps administrator use AWS CloudFormation to create a solution?

- A. Use Amazon EC2 user data in a CloudFormation template.
- B. Use nested stacks to provision resources.
- C. Use parameters in a CloudFormation template.
- D. Use stack policies to provision resources.

**Answer: C**

**Explanation:**

Reuse templates to replicate stacks in multiple environments. After you have your stacks and resources set up, you can reuse your templates to replicate your infrastructure in multiple environments. For example, you can create environments for development, testing, and production so that you can test changes before implementing them into production. To make templates reusable, use the parameters, mappings, and conditions sections so that you can customize your stacks when you create them. For example, for your development environments, you can specify a lower-cost instance type compared to your production environment, but all other configurations and settings remain the same. (<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html#reuse>)

**NEW QUESTION 10**

- (Exam Topic 1)

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

D. Create a Lambda function to restart the EC2 instance, triggered by EC2 health checks

**Answer: B**

#### NEW QUESTION 11

- (Exam Topic 1)

A company plans to run a public web application on Amazon EC2 instances behind an Elastic Load Balancer (ELB). The company's security team wants to protect the website by using AWS Certificate Manager (ACM) certificates. The ELB must automatically redirect any HTTP requests to HTTPS. Which solution will meet these requirements?

- A. Create an Application Load Balancer that has one HTTPS listener on port 80. Attach an SSL/TLS certificate to listener port 80. Create a rule to redirect requests from HTTP to HTTPS.
- B. Create an Application Load Balancer that has one HTTP listener on port 80 and one HTTPS protocol listener on port 443. Attach an SSL/TLS certificate to listener port 443. Create a rule to redirect requests from port 80 to port 443.
- C. Create an Application Load Balancer that has two TCP listeners on port 80 and port 443. Attach an SSL/TLS certificate to listener port 443. Create a rule to redirect requests from port 80 to port 443.
- D. Create a Network Load Balancer that has two TCP listeners on port 80 and port 443. Attach an SSL/TLS certificate to listener port 443. Create a rule to redirect requests from port 80 to port 443.

**Answer: B**

#### NEW QUESTION 13

- (Exam Topic 1)

A company has an organization in AWS Organizations. The company uses shared VPCs to provide networking resources across accounts. A SysOps administrator has been able to successfully launch and manage Amazon EC2 instances in a participant account. However, the SysOps administrator is now receiving an InstanceLimitExceeded error when the SysOps administrator tries to launch a new EC2 instance. What should the SysOps administrator do to resolve this error?

- A. Request an instance quota increase from the account that owns the VPC.
- B. Launch additional EC2 instances in a different AWS Region.
- C. Request an instance quota increase from the participant account.
- D. Launch additional EC2 instances by using a different Amazon Machine Image (AMI).

**Answer: A**

#### NEW QUESTION 16

- (Exam Topic 1)

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

**Answer: B**

#### NEW QUESTION 20

- (Exam Topic 1)

A company uses AWS Organizations. A SysOps administrator wants to use AWS Compute Optimizer and AWS tag policies in the management account to govern all member accounts in the billing family. The SysOps administrator navigates to the AWS Organizations console but cannot activate tag policies through the management account. What could be the reason for this issue?

- A. All features have not been enabled in the organization.
- B. Consolidated billing has not been enabled.
- C. The member accounts do not have tags enabled for cost allocation.
- D. The member accounts have not manually enabled trusted access for Compute Optimizer.

**Answer: C**

#### NEW QUESTION 24

- (Exam Topic 1)

A company stores files on 50 Amazon S3 buckets in the same AWS Region. The company wants to connect to the S3 buckets securely over a private connection from its Amazon EC2 instances. The company needs a solution that produces no additional cost. Which solution will meet these requirements?

- A. Create a gateway VPC endpoint for each S3 bucket.
- B. Attach the gateway VPC endpoints to each subnet inside the VPC.
- C. Create an interface VPC endpoint for each S3 bucket.
- D. Attach the interface VPC endpoints to each subnet inside the VPC.
- E. Create one gateway VPC endpoint for all the S3 buckets.
- F. Add the gateway VPC endpoint to the VPC route table.
- G. Create one interface VPC endpoint for all the S3 buckets.
- H. Add the interface VPC endpoint to the VPC route table.

**Answer: C**

**NEW QUESTION 25**

- (Exam Topic 1)

A company hosts a web portal on Amazon EC2 instances. The web portal uses an Elastic Load Balancer (ELB) and Amazon Route 53 for its public DNS service. The ELB and the EC2 instances are deployed by way of a single AWS CloudFormation stack in the us-east-1 Region. The web portal must be highly available across multiple Regions.

Which configuration will meet these requirements?

- A. Deploy a copy of the stack in the us-west-2 Region
- B. Create a single start of authority (SOA) record in Route 53 that includes the IP address from each EL
- C. Configure the SOA record with health check
- D. Use the ELB in us-east-1 as the primary record and the ELB in us-west-2 as the secondary record.
- E. Deploy a copy of the stack in the us-west-2 Region
- F. Create an additional A record in Route 53 that includes the ELB in us-west-2 as an alias target
- G. Configure the A records with a failover routing policy and health check
- H. Use the ELB in us-east-1 as the primary record and the ELB in us-west-2 as the secondary record.
- I. Deploy a new group of EC2 instances in the us-west-2 Region
- J. Associate the new EC2 instances with the existing ELB, and configure load balancer health checks on all EC2 instance
- K. Configure the ELB to update Route 53 when EC2 instances in us-west-2 fail health checks.
- L. Deploy a new group of EC2 instances in the us-west-2 Region
- M. Configure EC2 health checks on all EC2 instances in each Region
- N. Configure a peering connection between the VPC
- O. Use the VPC in us-east-1 as the primary record and the VPC in us-west-2 as the secondary record.

**Answer: B**

**Explanation:**

When you create a hosted zone, Route 53 automatically creates a name server (NS) record and a start of authority (SOA) record for the zone.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/migrate-dns-domain-in-use.html#migrate-dns-crea>

[https://en.wikipedia.org/wiki/SOA\\_record](https://en.wikipedia.org/wiki/SOA_record)

**NEW QUESTION 27**

- (Exam Topic 1)

A company runs an application on an Amazon EC2 instance. A SysOps administrator creates an Auto Scaling group and an Application Load Balancer (ALB) to handle an increase in demand. However, the EC2 instances are failing their health check.

What should the SysOps administrator do to troubleshoot this issue?

- A. Verify that the Auto Scaling group is configured to use all AWS Regions.
- B. Verify that the application is running on the protocol and the port that the listener is expecting.
- C. Verify the listener priority in the ALB. Change the priority if necessary.
- D. Verify the maximum number of instances in the Auto Scaling group. Change the number if necessary.

**Answer: B**

**NEW QUESTION 28**

- (Exam Topic 1)

A SysOps administrator is optimizing the cost of a workload. The workload is running in multiple AWS Regions and is using AWS Lambda with Amazon EC2 On-Demand Instances for the compute. The overall usage is predictable. The amount of compute that is consumed in each Region varies, depending on the users' locations.

Which approach should the SysOps administrator use to optimize this workload?

- A. Purchase Compute Savings Plans based on the usage during the past 30 days
- B. Purchase Convertible Reserved Instances by calculating the usage baseline.
- C. Purchase EC2 Instance Savings Plan based on the usage during the past 30 days
- D. Purchase Standard Reserved Instances by calculating the usage baseline.

**Answer: C**

**NEW QUESTION 30**

- (Exam Topic 1)

A company hosts a web application on an Amazon EC2 instance in a production VPC. Client connections to the application are failing. A SysOps administrator inspects the VPC flow logs and finds the following entry:

```
2 111122223333 eni-####> 192.0.2.15 203.0.113.56 40711 443 6 1 40 1418530010 1418530070 REJECT OK
```

What is a possible cause of these failed connections?

- A. A security group is denying traffic on port 443.
- B. The EC2 instance is shut down.
- C. The network ACL is blocking HTTPS traffic.
- D. The VPC has no internet gateway attached.

**Answer: A**

**Explanation:**

<https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs-records-examples.html#flow-log-example-accepted>

<https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs-records-examples.html#>

Accepted and rejected traffic: In this example, RDP traffic (destination port 3389, TCP protocol) to network interface eni-1235b8ca123456789 in account 123456789010 was rejected. 2 123456789010

```
eni-1235b8ca123456789 172.31.9.69 172.31.9.12 49761 3389 6 20 4249 1418530010 1418530070 REJECT OK
```

### NEW QUESTION 35

- (Exam Topic 1)

A SysOps administrator needs to design a high-traffic static website. The website must be highly available and must provide the lowest possible latency to users across the globe.

Which solution will meet these requirements?

- A. Create an Amazon S3 bucket, and upload the website content to the S3 bucket
- B. Create an Amazon CloudFront distribution in each AWS Region, and set the S3 bucket as the origin
- C. Use Amazon Route 53 to create a DNS record that uses a geolocation routing policy to route traffic to the correct CloudFront distribution based on where the request originates.
- D. Create an Amazon S3 bucket, and upload the website content to the S3 bucket
- E. Create an Amazon CloudFront distribution, and set the S3 bucket as the origin
- F. Use Amazon Route 53 to create an alias record that points to the CloudFront distribution.
- G. Create an Application Load Balancer (ALB) and a target group
- H. Create an Amazon EC2 Auto Scaling group with at least two EC2 instances in the associated target group
- I. Store the website content on the EC2 instance
- J. Use Amazon Route 53 to create an alias record that points to the ALB.
- K. Create an Application Load Balancer (ALB) and a target group in two Region
- L. Create an Amazon EC2 Auto Scaling group in each Region with at least two EC2 instances in each target group
- M. Store the website content on the EC2 instance
- N. Use Amazon Route 53 to create a DNS record that uses a geolocation routing policy to route traffic to the correct ALB based on where the request originates.

**Answer: B**

### NEW QUESTION 37

- (Exam Topic 1)

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? (Select 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 port used by the application developer does not match the port specified in the RDS configuration.
- E. The database is still being created and is not available for connectivity.

**Answer: CD**

### NEW QUESTION 41

- (Exam Topic 1)

A company has an initiative to reduce costs associated with Amazon EC2 and AWS Lambda. Which action should a SysOps administrator take to meet these requirements?

- A. Analyze the AWS Cost and Usage Report by using Amazon Athena to identify cost savings.
- B. Create an AWS Budgets alert to alarm when account spend reaches 80% of the budget.
- C. Purchase Reserved Instances through the Amazon EC2 console.
- D. Use AWS Compute Optimizer and take action on the provided recommendations.

**Answer: D**

### NEW QUESTION 42

- (Exam Topic 1)

A company has a stateful web application that is hosted on Amazon EC2 instances in an Auto Scaling group. The instances run behind an Application Load Balancer (ALB) that has a single target group. The ALB is configured as the origin in an Amazon CloudFront distribution. Users are reporting random logouts from the web application.

Which combination of actions should a SysOps administrator take to resolve this problem? (Select TWO.)

- A. Change to the least outstanding requests algorithm on the ALB target group.
- B. Configure cookie forwarding in the CloudFront distribution cache behavior.
- C. Configure header forwarding in the CloudFront distribution cache behavior.
- D. Enable group-level stickiness on the ALB listener rule.
- E. Enable sticky sessions on the ALB target group.

**Answer: BE**

#### Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Cookies.html>

You can configure each cache behavior to do one of the following: Forward all cookies to your origin – CloudFront includes all cookies sent by the viewer when it forwards requests to the origin. <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/sticky-sessions.html>

By default, an Application Load Balancer routes each request independently to a registered target based on the chosen load-balancing algorithm.

### NEW QUESTION 46

- (Exam Topic 1)

A company is running a serverless application on AWS Lambda. The application stores data in an Amazon RDS for MySQL DB instance. Usage has steadily increased and recently there have been numerous "too many connections" errors when the Lambda function attempts to connect to the database. The company already has configured the database to use the maximum `max_connections` value that is possible.

What should a SysOps administrator do to resolve these errors?

- A. Create a read replica of the database Use Amazon Route 53 to create a weighted DNS record that contains both databases
- B. Use Amazon RDS Proxy to create a proxy Update the connection string in the Lambda function
- C. Increase the value in the max\_connect\_errors parameter in the parameter group that the database uses
- D. Update the Lambda function's reserved concurrency to a higher value

**Answer: B**

**Explanation:**

<https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/>

RDS Proxy acts as an intermediary between your application and an RDS database. RDS Proxy establishes and manages the necessary connection pools to your database so that your application creates fewer database connections. Your Lambda functions interact with RDS Proxy instead of your database instance. It handles the connection pooling necessary for scaling many simultaneous connections created by concurrent Lambda functions. This allows your Lambda applications to reuse existing connections, rather than creating new connections for every function invocation.

Check "Database proxy for Amazon RDS" section in the link to see how RDS proxy help Lambda handle huge connections to RDS MySQL

<https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/>

**NEW QUESTION 47**

- (Exam Topic 1)

A company using AWS Organizations requires that no Amazon S3 buckets in its production accounts should ever be deleted.

What is the SIMPLEST approach the SysOps administrator can take to ensure S3 buckets in those accounts can never be deleted?

- A. Set up MFA Delete on all the S3 buckets to prevent the buckets from being deleted.
- B. Use service control policies to deny the s3:DeleteBucket action on all buckets in production accounts.
- C. Create an IAM group that has an IAM policy to deny the s3:DeleteBucket action on all buckets in production accounts.
- D. Use AWS Shield to deny the s3:DeleteBucket action on the AWS account instead of all S3 buckets.

**Answer: B**

**Explanation:**

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scps.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html)

If you're using AWS Organizations, check the service control policies for any statements that explicitly deny Amazon S3 access. In particular, check the service control policies for statements denying the s3:PutBucketPolicy action.

<https://aws.amazon.com/tw/premiumsupport/knowledge-center/s3-access-denied-bucket-policy/>

**NEW QUESTION 49**

- (Exam Topic 1)

A company uses an Amazon Elastic File System (Amazon EFS) file system to share files across many Linux Amazon EC2 instances. A SysOps administrator notices that the file system's PercentIOLimit metric is consistently at 100% for 15 minutes or longer. The SysOps administrator also notices that the application that reads and writes to that file system is performing poorly. They application requires high throughput and IOPS while accessing the file system.

What should the SysOps administrator do to remediate the consistently high PercentIOLimit metric?

- A. Create a new EFS file system that uses Max I/O performance mod
- B. Use AWS DataSync to migrate data to the new EFS file system.
- C. Create an EFS lifecycle policy to transition future files to the Infrequent Access (IA) storage class to improve performanc
- D. Use AWS DataSync to migrate existing data to IA storage.
- E. Modify the existing EFS file system and activate Max I/O performance mode.
- F. Modify the existing EFS file system and activate Provisioned Throughput mode.

**Answer: A**

**Explanation:**

To support a wide variety of cloud storage workloads, Amazon EFS offers two performance modes, General Purpose mode and Max I/O mode. You choose a file system's performance mode when you create it, and it cannot be changed. If the PercentIOLimit percentage returned was at or near 100 percent for a significant amount of time during the test, your application should use the Max I/O performance mode. <https://docs.aws.amazon.com/efs/latest/ug/performance.html>

**NEW QUESTION 51**

- (Exam Topic 1)

A company wants to be alerted through email when IAM CreateUser API calls are made within its AWS account.

Which combination of actions should a SysOps administrator take to meet this requirement? (Choose two.)

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with AWS CloudTrail as the event source and IAM CreateUser as the specific API call for the event pattern.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with Amazon CloudSearch as the event source and IAM CreateUser as the specific API call for the event pattern.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with AWS IAM Access Analyzer as the event source and IAM CreateUser as the specific API call for the event pattern.
- D. Use an Amazon Simple Notification Service (Amazon SNS) topic as an event target with an email subscription.
- E. Use an Amazon Simple Email Service (Amazon SES) notification as an event target with an email subscription.

**Answer: AD**

**Explanation:**

<https://aws.amazon.com/blogs/security/how-to-receive-alerts-when-your-iam-configuration-changes/>

**NEW QUESTION 54**

- (Exam Topic 1)

A SysOps administrator creates two VPCs, VPC1 and VPC2, in a company's AWS account The SysOps administrator deploys a Linux Amazon EC2 instance in VPC1 and deploys an Amazon RDS for MySQL DB instance in VPC2. The DB instance is deployed in a private subnet. An application that runs on the EC2 instance needs to connect to the database.

What should the SysOps administrator do to give the EC2 instance the ability to connect to the database?

- A. Enter the DB instance connection string into the VPC1 route table.
- B. Configure VPC peering between the two VPCs.
- C. Add the same IPv4 CIDR range for both VPCs.
- D. Connect to the DB instance by using the DB instance's public IP address.

**Answer: B**

**Explanation:**

VPC peering allows two VPCs to communicate with each other securely. By configuring VPC peering between the two VPCs, the SysOps administrator will be able to give the EC2 instance in VPC1 the ability to connect to the database in VPC2. Once the VPC peering is configured, the EC2 instance will be able to communicate with the database using the private IP address of the DB instance in the private subnet.

**NEW QUESTION 58**

- (Exam Topic 1)

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 --na flag.

**Answer: C**

**NEW QUESTION 62**

- (Exam Topic 1)

A company monitors its account activity using AWS CloudTrail, and is concerned that some log files are being tampered with after the logs have been delivered to the account's Amazon S3 bucket.

Moving forward, how can the SysOps administrator confirm that the log files have not been modified after being delivered to the S3 bucket?

- A. Stream the CloudTrail logs to Amazon CloudWatch Logs to store logs at a secondary location.
- B. Enable log file integrity validation and use digest files to verify the hash value of the log file.
- C. Replicate the S3 log bucket across regions, and encrypt log files with S3 managed keys.
- D. Enable S3 server access logging to track requests made to the log bucket for security audits.

**Answer: B**

**Explanation:**

When you enable log file integrity validation, CloudTrail creates a hash for every log file that it delivers. Every hour, CloudTrail also creates and delivers a file that references the log files for the last hour and contains a hash of each. This file is called a digest file. CloudTrail signs each digest file using the private key of a public and private key pair. After delivery, you can use the public key to validate the digest file. CloudTrail uses different key pairs for each AWS region  
<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

**NEW QUESTION 64**

- (Exam Topic 1)

A SysOps administrator is required to monitor free space on Amazon EBS volumes attached to Microsoft Windows-based Amazon EC2 instances within a company's account. The administrator must be alerted to potential issues.

What should the administrator do to receive email alerts before low storage space affects EC2 instance performance?

- A. Use built-in Amazon CloudWatch metrics, and configure CloudWatch alarms and an Amazon SNS topic for email notifications
- B. Use AWS CloudTrail logs and configure the trail to send notifications to an Amazon SNS topic.
- C. Use the Amazon CloudWatch agent to send disk space metrics, then set up CloudWatch alarms using an Amazon SNS topic.
- D. Use AWS Trusted Advisor and enable email notification alerts for EC2 disk space

**Answer: C**

**NEW QUESTION 68**

- (Exam Topic 1)

A company updates its security policy to clarify cloud hosting arrangements for regulated workloads. Workloads that are identified as sensitive must run on hardware that is not shared with other customers or with other AWS accounts within the company.

Which solution will ensure compliance with this policy?

- A. Deploy workloads only to Dedicated Hosts.
- B. Deploy workloads only to Dedicated Instances.
- C. Deploy workloads only to Reserved Instances.
- D. Place all instances in a dedicated placement group.

**Answer: A**

**Explanation:**

Dedicated Hosts are physical servers that are dedicated to a single customer, ensuring that the customer's workloads are not shared with other customers or with other AWS accounts within the company. This will ensure that the company's security policy is followed and that sensitive workloads are running on hardware that is not shared with other customers or with other AWS accounts within the company.

**NEW QUESTION 70**

- (Exam Topic 1)

An environment consists of 100 Amazon EC2 Windows instances. The Amazon CloudWatch agent is deployed and running on all EC2 instances with a baseline configuration file to capture log files. There is a new requirement to capture the DHCP log files that exist on 50 of the instances. What is the MOST operational efficient way to meet this new requirement?

- A. Create an additional CloudWatch agent configuration file to capture the DHCP logs. Use the AWS Systems Manager Run Command to restart the CloudWatch agent on each EC2 instance with the append-config option to apply the additional configuration file.
- B. Log in to each EC2 instance with administrator rights. Create a PowerShell script to push the needed baseline log files and DHCP log files to CloudWatch.
- C. Run the CloudWatch agent configuration file wizard on each EC2 instance. Verify that the base log files are included and add the DHCP log files during the wizard creation process.
- D. Run the CloudWatch agent configuration file wizard on each EC2 instance and select the advanced detail level.
- E. This will capture the operating system log files.

**Answer: A**

**NEW QUESTION 72**

- (Exam Topic 1)

A company has attached the following policy to an IAM user:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "rds:Describe*",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    },
    {
      "Effect": "Deny",
      "NotAction": {
        "ec2:*"
      }
    },
    {
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "ec2:*",
        "s3:GetObject"
      ],
      "Resource": "*"
    }
  ]
}
```

Which of the following actions are allowed for the IAM user?

- A. Amazon RDS DescribeDBInstances action in the us-east-1 Region
- B. Amazon S3 PutObject operation in a bucket named testbucket
- C. Amazon EC2 Describe Instances action in the us-east-1 Region
- D. Amazon EC2 AttachNetworkInterface action in the eu-west-1 Region

**Answer:** C

**NEW QUESTION 77**

- (Exam Topic 1)

A company uses an Amazon Simple Queue Service (Amazon SQS) standard queue with its application. The application sends messages to the queue with unique message bodies. The company decides to switch to an SQS FIFO queue. What must the company do to migrate to an SQS FIFO queue?

- A. Create a new SQS FIFO queue. Turn on content-based deduplication on the new FIFO queue. Update the application to include a message group ID in the messages.
- B. Create a new SQS FIFO queue. Update the application to include the DelaySeconds parameter in the messages.
- C. Modify the queue type from SQS standard to SQS FIFO. Turn off content-based deduplication on the queue. Update the application to include a message group ID in the messages.
- D. Modify the queue type from SQS standard to SQS FIFO. Update the application to send messages with identical message bodies and to include the DelaySeconds parameter in the messages.

**Answer:** A

**Explanation:**

FIFO queues don't support per-message delays, only per-queue delays. If your application sets the same value of the DelaySeconds parameter on each message, you must modify your application to remove the per-message delay and set DelaySeconds on the entire queue instead.

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues-moving.html>

**NEW QUESTION 78**

- (Exam Topic 1)

A company is expanding its use of AWS services across its portfolios. The company wants to provision AWS accounts for each team to ensure a separation of business processes for security compliance and billing. Account creation and bootstrapping should be completed in a scalable and efficient way so new accounts are created with a defined baseline and governance guardrails in place. A SysOps administrator needs to design a provisioning process that saves time and resources.

Which action should be taken to meet these requirements?

- A. Automate using AWS Elastic Beanstalk to provision the AWS accounts, set up infrastructure, and integrate with AWS Organizations.
- B. Create bootstrapping scripts in AWS OpsWorks and combine them with AWS CloudFormation templates to provision accounts and infrastructure.
- C. Use AWS Config to provision accounts and deploy instances using AWS Service Catalog.
- D. Use AWS Control Tower to create a template in Account Factory and use the template to provision new accounts.

**Answer:** D

**NEW QUESTION 83**

- (Exam Topic 1)

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

**Answer:** D

**Explanation:**

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingEncryption.html>

**NEW QUESTION 88**

- (Exam Topic 1)

A SysOps administrator is evaluating Amazon Route 53 DNS options to address concerns about high availability for an on-premises website. The website consists of two servers: a primary active server and a secondary passive server. Route 53 should route traffic to the primary server if the associated health check returns 2xx or 3xx HTTP codes. All other traffic should be directed to the secondary passive server. The failover record type, set ID, and routing policy have been set appropriately for both primary and secondary servers.

Which next step should be taken to configure Route 53?

- A. Create an A record for each server.
- B. Associate the records with the Route 53 HTTP health check.
- C. Create an A record for each server.
- D. Associate the records with the Route 53 TCP health check.
- E. Create an alias record for each server with evaluate target health set to yes.
- F. Associate the records with the Route 53 HTTP health check.
- G. Create an alias record for each server with evaluate target health set to yes.
- H. Associate the records with the Route 53 TCP health check.

**Answer:** A

**NEW QUESTION 91**

- (Exam Topic 1)

A SysOps administrator configures an Amazon S3 gateway endpoint in a VPC. The private subnets inside the VPC do not have outbound internet access. A user

logs in to an Amazon EC2 instance in one of the private subnets and cannot upload a file to an Amazon S3 bucket in the same AWS Region. Which solution will solve this problem?

- A. Update the EC2 instance role policy to allow s3:PutObject access to the target S3 bucket.
- B. Update the EC2 security group to allow outbound traffic to 0.0.0.0/0 for port 80.
- C. Update the EC2 subnet route table to include the S3 prefix list destination routes to the S3 gateway endpoint.
- D. Update the S3 bucket policy to allow s3:PutObject access from the private subnet CIDR block.

**Answer: C**

#### NEW QUESTION 94

- (Exam Topic 1)

A company has a high-performance Windows workload. The workload requires a storage volume that provides consistent performance of 10,000 IOPS. The company does not want to pay for additional unneeded capacity to achieve this performance. Which solution will meet these requirements with the LEAST cost?

- A. Use a Provisioned IOPS SSD (io1) Amazon Elastic Block Store (Amazon EBS) volume that is configured with 10,000 provisioned IOPS.
- B. Use a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume that is configured with 10,000 provisioned IOPS.
- C. Use an Amazon Elastic File System (Amazon EFS) file system with Max I/O mode.
- D. Use an Amazon FSx for Windows File Server file system that is configured with 10,000 IOPS.

**Answer: A**

#### NEW QUESTION 96

- (Exam Topic 1)

A company is expanding globally and needs to back up data on Amazon Elastic Block Store (Amazon EBS) volumes to a different AWS Region. Most of the EBS volumes that store the data are encrypted, but some of the EBS volumes are unencrypted. The company needs the backup data from all the EBS volumes to be encrypted.

Which solution will meet these requirements with the LEAST management overhead?

- A. Configure a lifecycle policy in Amazon Data Lifecycle Manager (Amazon DLM) to create the EBS volume snapshots with cross-Region backups enabled.
- B. Encrypt the snapshot copies by using AWS Key Management Service (AWS KMS).
- C. Create a point-in-time snapshot of the EBS volume.
- D. When the snapshot status is COMPLETED, copy the snapshots to another Region and set the Encrypted parameter to False.
- E. Create a point-in-time snapshot of the EBS volume.
- F. Copy the snapshots to an Amazon S3 bucket that uses server-side encryption.
- G. Turn on S3 Cross-Region Replication on the S3 bucket.
- H. Schedule an AWS Lambda function with the Python runtime.
- I. Configure the Lambda function to create the EBS volume snapshots, encrypt the unencrypted snapshots, and copy the snapshots to another Region.

**Answer: A**

#### Explanation:

Encrypt the snapshot copies by using AWS Key Management Service (AWS KMS). This solution will allow the company to automatically create encrypted snapshots of the EBS volumes and copy them to different AWS Regions with minimal effort.

#### NEW QUESTION 101

- (Exam Topic 1)

A company's SysOps administrator deploys four new Amazon EC2 instances by using the standard Amazon Linux 2 Amazon Machine Image (AMI). The company needs to be able to use AWS Systems Manager to manage the instances. The SysOps administrator notices that the instances do not appear in the Systems Manager console.

What must the SysOps administrator do to resolve this issue?

- A. Connect to each instance by using SSH. Install Systems Manager Agent on each instance. Configure Systems Manager Agent to start automatically when the instances start up.
- B. Use AWS Certificate Manager (ACM) to create a TLS certificate. Import the certificate into each instance. Configure Systems Manager Agent to use the TLS certificate for secure communications.
- C. Connect to each instance by using SSH. Create an ssm-user account. Add the ssm-user account to the /etc/sudoers.d directory.
- D. Attach an IAM instance profile to the instances. Ensure that the instance profile contains the AmazonSSMManagedInstanceCore policy.

**Answer: D**

#### NEW QUESTION 104

- (Exam Topic 1)

A company uses an Amazon CloudFront distribution to deliver its website. Traffic logs for the website must be centrally stored and all data must be encrypted at rest.

Which solution will meet these requirements?

- A. Create an Amazon OpenSearch Service (Amazon Elasticsearch Service) domain with internet access and server-side encryption that uses the default AWS managed key. Configure CloudFront to use the Amazon OpenSearch Service (Amazon Elasticsearch Service) domain as a log destination.
- B. Create an Amazon OpenSearch Service (Amazon Elasticsearch Service) domain with VPC access and server-side encryption that uses AES-256. Configure CloudFront to use the Amazon OpenSearch Service (Amazon Elasticsearch Service) domain as a log destination.
- C. Create an Amazon S3 bucket that is configured with default server-side encryption that uses AES-256. Configure CloudFront to use the S3 bucket as a log destination.
- D. Create an Amazon S3 bucket that is configured with no default encryption. Enable encryption in the CloudFront distribution and use the S3 bucket as a log destination.

**Answer: C**

### NEW QUESTION 108

- (Exam Topic 1)

A company has an Amazon CloudFront distribution that uses an Amazon S3 bucket as its origin. During a review of the access logs, the company determines that some requests are going directly to the S3 bucket by using the website hosting endpoint. A SysOps administrator must secure the S3 bucket to allow requests only from CloudFront.

What should the SysOps administrator do to meet this requirement?

- A. Create an origin access identity (OAI) in CloudFront
- B. Associate the OAI with the distribution
- C. Remove access to and from other principals in the S3 bucket policy
- D. Update the S3 bucket policy to allow access only from the OAI.
- E. Create an origin access identity (OAI) in CloudFront
- F. Associate the OAI with the distribution
- G. Update the S3 bucket policy to allow access only from the OAI
- H. Create a new origin, and specify the S3 bucket as the new origin
- I. Update the distribution behavior to use the new origin
- J. Remove the existing origin.
- K. Create an origin access identity (OAI) in CloudFront
- L. Associate the OAI with the distribution
- M. Update the S3 bucket policy to allow access only from the OAI
- N. Disable website hosting
- O. Create a new origin, and specify the S3 bucket as the new origin
- P. Update the distribution behavior to use the new origin
- Q. Remove the existing origin.
- R. Update the S3 bucket policy to allow access only from the CloudFront distribution
- S. Remove access to and from other principals in the S3 bucket policy
- T. Disable website hosting
- . Create a new origin, and specify the S3 bucket as the new origin
- . Update the distribution behavior to use the new origin
- . Remove the existing origin.

**Answer: A**

### NEW QUESTION 109

- (Exam Topic 1)

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

**Answer: D**

### NEW QUESTION 113

- (Exam Topic 1)

A SysOps administrator is configuring an application on Amazon EC2 instances for a company. Teams in other countries will use the application over the internet. The company requires the application endpoint to have a static public IP address.

How should the SysOps administrator deploy the application to meet this requirement?

- A. Behind an Amazon API Gateway API
- B. Behind an Application Load Balancer
- C. Behind an internet-facing Network Load Balancer
- D. In an Amazon CloudFront distribution

**Answer: C**

### NEW QUESTION 115

- (Exam Topic 1)

An application accesses data through a file system interface. The application runs on Amazon EC2 instances in multiple Availability Zones, all of which must share the same data. While the amount of data is currently small, the company anticipates that it will grow to tens of terabytes over the lifetime of the application.

What is the MOST scalable storage solution to fulfill this requirement?

- A. Connect a large Amazon EBS volume to multiple instances and schedule snapshots.
- B. Deploy Amazon EFS in the VPC and create mount targets in multiple subnets.
- C. Launch an EC2 instance and share data using SMB/CIFS or NFS.
- D. Deploy an AWS Storage Gateway cached volume on Amazon EC2.

**Answer: B**

### NEW QUESTION 117

- (Exam Topic 1)

A company runs a web application on three Amazon EC2 instances behind an Application Load Balancer (ALB). The company notices that random periods of increased traffic cause a degradation in the application's performance. A SysOps administrator must scale the application to meet the increased traffic. Which solution meets these requirements?

- A. Create an Amazon CloudWatch alarm to monitor application latency and increase the size of each EC2 instance if the desired threshold is reached.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor application latency and add an EC2 instance to the ALB if the desired threshold is

reached.

- C. Deploy the application to an Auto Scaling group of EC2 instances with a target tracking scaling policy. Attach the ALB to the Auto Scaling group.
- D. Deploy the application to an Auto Scaling group of EC2 instances with a scheduled scaling policy. Attach the ALB to the Auto Scaling group.

**Answer: C**

**NEW QUESTION 118**

- (Exam Topic 1)

A large company is using AWS Organizations to manage hundreds of AWS accounts across multiple AWS Regions. The company has turned on AWS Config throughout the organization.

The company requires all Amazon S3 buckets to block public read access. A SysOps administrator must generate a monthly report that shows all the S3 buckets and whether they comply with this requirement.

Which combination of steps should the SysOps administrator take to collect this data? (Select TWO).

- A. Create an AWS Config aggregator in an aggregator account
- B. Use the organization as the source. Retrieve the compliance data from the aggregator.
- C. Create an AWS Config aggregator in each account
- D. Use an S3 bucket in an aggregator account as the destination
- E. Retrieve the compliance data from the S3 bucket
- F. Edit the AWS Config policy in AWS Organization
- G. Use the organization's management account to turn on the s3-bucket-public-read-prohibited rule for the entire organization.
- H. Use the AWS Config compliance report from the organization's management account
- I. Filter the results by resource, and select Amazon S3.
- J. Use the AWS Config API to apply the s3-bucket-public-read-prohibited rule in all accounts for all available Regions.

**Answer: CD**

**NEW QUESTION 122**

- (Exam Topic 1)

A company is using Amazon Elastic File System (Amazon EFS) to share a file system among several Amazon EC2 instances. As usage increases, users report that file retrieval from the EFS file system is slower than normal.

Which action should a SysOps administrator take to improve the performance of the file system?

- A. Configure the file system for Provisioned Throughput.
- B. Enable encryption in transit on the file system.
- C. Identify any unused files in the file system, and remove the unused files.
- D. Resize the Amazon Elastic Block Store (Amazon EBS) volume of each of the EC2 instances.

**Answer: A**

**NEW QUESTION 127**

- (Exam Topic 1)

A company is storing backups in an Amazon S3 bucket. The backups must not be deleted for at least 3 months after the backups are created.

What should a SysOps administrator do to meet this requirement?

- A. Configure an IAM policy that denies the s3:DeleteObject action for all user
- B. Three months after an object is written, remove the policy.
- C. Enable S3 Object Lock on a new S3 bucket in compliance mode
- D. Place all backups in the new S3 bucket with a retention period of 3 months.
- E. Enable S3 Versioning on the existing S3 bucket
- F. Configure S3 Lifecycle rules to protect the backups.
- G. Enable S3 Object Lock on a new S3 bucket in governance mode
- H. Place all backups in the new S3 bucket with a retention period of 3 months.

**Answer: D**

**Explanation:**

To meet the requirements of the workload, a SysOps administrator should enable S3 Object Lock on a new S3 bucket in governance mode and place all backups in the new S3 bucket with a retention period of 3 months.

This will ensure that the backups are not deleted for at least 3 months after they are created. The other solutions (configuring an IAM policy that denies the s3:DeleteObject action for all users, enabling S3 Object Lock on a new S3 bucket in compliance mode, or enabling S3 Versioning on the existing S3 bucket and configuring S3 Lifecycle rules to protect the backups) will not meet the requirements, as they do not provide a way to ensure that the backups are not deleted for at least 3 months after they are created.

**NEW QUESTION 131**

- (Exam Topic 1)

Application A runs on Amazon EC2 instances behind a Network Load Balancer (NLB). The EC2 instances are in an Auto Scaling group and are in the same subnet that is associated with the NLB. Other applications from an on-premises environment cannot communicate with Application A on port 8080.

To troubleshoot the issue, a SysOps administrator analyzes the flow logs. The flow logs include the following records:

```
2 123456789010 eni-1235b8ca123456789 192.168.0.13 172.31.16.139 59003 8080 1 4 336 1432917027 1432917142 ACCEPT OK
2 123456789010 eni-1235b8ca123456789 172.31.16.139 192.168.0.13 8080 59003 1 4 336 1432917094 1432917142 REJECT OK
```

What is the reason for the rejected traffic?

- A. The security group of the EC2 instances has no Allow rule for the traffic from the NLB.
- B. The security group of the NLB has no Allow rule for the traffic from the on-premises environment.
- C. The ACL of the on-premises environment does not allow traffic to the AWS environment.
- D. The network ACL that is associated with the subnet does not allow outbound traffic for the ephemeral port range.

**Answer:** A

**NEW QUESTION 134**

- (Exam Topic 1)

A global company handles a large amount of personally identifiable information (PII) through an internal web portal. The company's application runs in a corporate data center that is connected to AWS through an AWS Direct Connect connection. The application stores the PII in Amazon S3. According to a compliance requirement, traffic from the web portal to Amazon S3 must not travel across the internet.

What should a SysOps administrator do to meet the compliance requirement?

- A. Provision an interface VPC endpoint for Amazon S3. Modify the application to use the interface endpoint.
- B. Configure AWS Network Firewall to redirect traffic to the internal S3 address.
- C. Modify the application to use the S3 path-style endpoint.
- D. Set up a range of VPC network ACLs to redirect traffic to the Internal S3 address.

**Answer:** B

**NEW QUESTION 135**

- (Exam Topic 1)

A company has a stateless application that runs on four Amazon EC2 instances. The application requires four instances at all times to support all traffic. A SysOps administrator must design a highly available, fault-tolerant architecture that continually supports all traffic if one Availability Zone becomes unavailable.

Which configuration meets these requirements?

- A. Deploy two Auto Scaling groups in two Availability Zones with a minimum capacity of two instances in each group.
- B. Deploy an Auto Scaling group across two Availability Zones with a minimum capacity of four instances.
- C. Deploy an Auto Scaling group across three Availability Zones with a minimum capacity of four instances.
- D. Deploy an Auto Scaling group across three Availability Zones with a minimum capacity of six instances.

**Answer:** C

**NEW QUESTION 137**

- (Exam Topic 1)

A company stores sensitive data in an Amazon S3 bucket. The company must log all access attempts to the S3 bucket. The company's risk team must receive immediate notification about any delete events.

Which solution will meet these requirements?

- A. Enable S3 server access logging for audit log
- B. Set up an Amazon Simple Notification Service (Amazon SNS) notification for the S3 bucket
- C. Select DeleteObject for the event type for the alert system.
- D. Enable S3 server access logging for audit log
- E. Launch an Amazon EC2 instance for the alert system. Run a cron job on the EC2 instance to download the access logs each day and to scan for a DeleteObject event.
- F. Use Amazon CloudWatch Logs for audit log
- G. Use Amazon CloudWatch alarms with an Amazon Simple Notification Service (Amazon SNS) notification for the alert system.
- H. Use Amazon CloudWatch Logs for audit log
- I. Launch an Amazon EC2 instance for The alert system. Run a cron job on the EC2 Instance each day to compare the list of the items with the list from the previous day
- J. Configure the cron job to send a notification if an item is missing.

**Answer:** A

**Explanation:**

To meet the requirements of logging all access attempts to the S3 bucket and receiving immediate notification about any delete events, the company can enable S3 server access logging and set up an Amazon Simple Notification Service (Amazon SNS) notification for the S3 bucket. The S3 server access logs will record all access attempts to the bucket, including delete events, and the SNS notification can be configured to send an alert when a DeleteObject event occurs.

**NEW QUESTION 139**

- (Exam Topic 1)

A large company is using AWS Organizations to manage its multi-account AWS environment. According to company policy, all users should have read-level access to a particular Amazon S3 bucket in a central account. The S3 bucket data should not be available outside the organization. A SysOps administrator must set up the permissions and add a bucket policy to the S3 bucket.

Which parameters should be specified to accomplish this in the MOST efficient manner?

- A. Specify '\*' as the principal and PrincipalOrgId as a condition.
- B. Specify all account numbers as the principal.
- C. Specify PrincipalOrgId as the principal.
- D. Specify the organization's management account as the principal.

**Answer:** C

**NEW QUESTION 140**

- (Exam Topic 1)

A company is using Amazon Elastic Container Service (Amazon ECS) to run a containerized application on Amazon EC2 instances. A SysOps administrator needs to monitor only traffic flows between the ECS tasks.

Which combination of steps should the SysOps administrator take to meet this requirement? (Select TWO.)

- A. Configure Amazon CloudWatch Logs on the elastic network interface of each task.
- B. Configure VPC Flow Logs on the elastic network interface of each task.

- C. Specify the awsvpc network mode in the task definition.
- D. Specify the bridge network mode in the task definition.
- E. Specify the host network mode in the task definition.

**Answer:** AE

#### NEW QUESTION 143

- (Exam Topic 1)

A company is undergoing an external audit of its systems, which run wholly on AWS. A SysOps administrator must supply documentation of Payment Card Industry Data Security Standard (PCI DSS) compliance for the infrastructure managed by AWS. Which set of action should the SysOps administrator take to meet this requirement?

- A. Download the applicable reports from the AWS Artifact portal and supply these to the auditors.
- B. Download complete copies of the AWS CloudTrail log files and supply these to the auditors.
- C. Download complete copies of the AWS CloudWatch logs and supply these to the auditors.
- D. Provide the auditors with administrative access to the production AWS account so that the auditors can determine compliance.

**Answer:** A

#### NEW QUESTION 145

- (Exam Topic 1)

A company has a critical serverless application that uses multiple AWS Lambda functions. Each Lambda function generates 1 GB of log data daily in its own Amazon CloudWatch Logs log group. The company's security team asks for a count of application errors, grouped by type, across all of the log groups. What should a SysOps administrator do to meet this requirement?

- A. Perform a CloudWatch Logs Insights query that uses the stats command and count function.
- B. Perform a CloudWatch Logs search that uses the groupby keyword and count function.
- C. Perform an Amazon Athena query that uses the SELECT and GROUP BY keywords.
- D. Perform an Amazon RDS query that uses the SELECT and GROUP BY keywords.

**Answer:** A

#### NEW QUESTION 150

- (Exam Topic 1)

A company is implementing a monitoring solution that is based on machine learning. The monitoring solution consumes Amazon EventBridge (Amazon CloudWatch Events) events that are generated by Amazon EC2 Auto Scaling. The monitoring solution provides detection of anomalous behavior such as unanticipated scaling events and is configured as an EventBridge (CloudWatch Events) API destination. During initial testing, the company discovers that the monitoring solution is not receiving events. However, Amazon CloudWatch is showing that the EventBridge (CloudWatch Events) rule is being invoked. A SysOps administrator must implement a solution to retrieve client error details to help resolve this issue. Which solution will meet these requirements with the LEAST operational effort?

- A. Create an EventBridge (CloudWatch Events) archive for the event pattern to replay the event
- B. Increase the logging on the monitoring solution
- C. Use replay to invoke the monitoring solution
- D. Examine the error details.
- E. Add an Amazon Simple Queue Service (Amazon SQS) standard queue as a dead-letter queue for the target
- F. Process the messages in the dead-letter queue to retrieve error details.
- G. Create a second EventBridge (CloudWatch Events) rule for the same event pattern to target an AWS Lambda function
- H. Configure the Lambda function to invoke the monitoring solution and to record the results to Amazon CloudWatch Log
- I. Examine the errors in the logs.
- J. Configure the EventBridge (CloudWatch Events) rule to send error messages to an Amazon Simple Notification Service (Amazon SNS) topic.

**Answer:** A

#### Explanation:

"In EventBridge, you can create an archive of events so that you can easily replay them at a later time. For example, you might want to replay events to recover from errors or to validate new functionality in your application." <https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-archive.html>

#### NEW QUESTION 154

- (Exam Topic 1)

A company's SysOps administrator needs to change the AWS Support plan for one of the company's AWS accounts. The account has multi-factor authentication (MFA) activated, and the MFA device is lost. What should the SysOps administrator do to sign in?

- A. Sign in as a root user by using email and phone verification
- B. Set up a new MFA device
- C. Change the root user password.
- D. Sign in as an IAM user with administrator permission
- E. Resynchronize the MFA token by using the IAM console.
- F. Sign in as an IAM user with administrator permission
- G. Reset the MFA device for the root user by adding a new device.
- H. Use the forgot-password process to verify the email address
- I. Set up a new password and MFA device.

**Answer:** A

#### NEW QUESTION 157

- (Exam Topic 1)

A company is creating a new multi-account architecture. A Sysops administrator must implement a login solution to centrally manage user access and permissions across all AWS accounts. The solution must be integrated with AWS Organizations and must be connected to a third-party Security Assertion Markup Language (SAML) 2.0 identity provider (IdP).

What should the SysOps administrator do to meet these requirements?

- A. Configure an Amazon Cognito user pool
- B. Integrate the user pool with the third-party IdP.
- C. Enable and configure AWS Single Sign-On with the third-party IdP.
- D. Federate the third-party IdP with AWS Identity and Access Management (IAM) for each AWS account in the organization.
- E. Integrate the third-party IdP directly with AWS Organizations.

**Answer: A**

#### NEW QUESTION 160

- (Exam Topic 1)

A SysOps administrator is setting up a fleet of Amazon EC2 instances in an Auto Scaling group for an application. The fleet should have 50% CPU available at that times to accommodate bursts of traffic. The load will increase significantly between the hours of 09:00 and 17:00, 7 days a week

How should the SysOps administrator configure the scaling of the EC2 instances to meet these requirements?

- A. Create a target tracking scaling policy that runs when the CPU utilization is higher than 90%
- B. Create a target tracking scaling policy that runs when the CPU utilization is higher than 50%. Create a scheduled scaling policy that ensures that the fleet is available at 09:00 Create a second scheduled scaling policy that scales in the fleet at 17:00
- C. Set the Auto Scaling group to start with 2 instances by setting the desired instances maximum instances, and minimum instances to 2 Create a scheduled scaling policy that ensures that the fleet is available at 09:00
- D. Create a scheduled scaling policy that ensures that the fleet is available at 09.00. Create a second scheduled scaling policy that scales in the fleet at 17:00

**Answer: B**

#### NEW QUESTION 161

- (Exam Topic 1)

A company wants to archive sensitive data on Amazon S3 Glacier. The company's regulatory and compliance requirements do not allow any modifications to the data by any account.

Which solution meets these requirements?

- A. Attach a vault lock policy to an S3 Glacier vault that contains the archived data
- B. Use the lock ID to validate the vault lock policy after 24 hours.
- C. Attach a vault lock policy to an S3 Glacier vault that contains the archived data
- D. Use the lock ID to validate the vault lock policy within 24 hours.
- E. Configure S3 Object Lock in governance mode
- F. Upload all files after 24 hours.
- G. Configure S3 Object Lock in governance mode
- H. Upload all files within 24 hours.

**Answer: B**

#### NEW QUESTION 166

- (Exam Topic 1)

A SysOps administrator is designing a solution for an Amazon RDS for PostgreSQL DB instance. Database credentials must be stored and rotated monthly. The applications that connect to the DB instance send

write-intensive traffic with variable client connections that sometimes increase significantly in a short period of time.

Which solution should a SysOps administrator choose to meet these requirements?

- A. Configure AWS Key Management Service (AWS KMS) to automatically rotate the keys for the DB instance
- B. Use RDS Proxy to handle the increases in database connections.
- C. Configure AWS Key Management Service (AWS KMS) to automatically rotate the keys for the DB instance
- D. Use RDS read replicas to handle the increases in database connections.
- E. Configure AWS Secrets Manager to automatically rotate the credentials for the DB instance
- F. Use RDS Proxy to handle the increases in database connections.
- G. Configure AWS Secrets Manager to automatically rotate the credentials for the DB instance
- H. Use RDS read replicas to handle the increases in database connections.

**Answer: A**

#### NEW QUESTION 171

- (Exam Topic 1)

A SysOps administrator has enabled AWS CloudTrail in an AWS account If CloudTrail is disabled it must be re-enabled immediately What should the SysOps administrator do to meet these requirements WITHOUT writing custom code"

- A. Add the AWS account to AWS Organizations Enable CloudTrail in the management account
- B. Create an AWS Config rule that is invoked when CloudTrail configuration changes Apply the AWS-ConfigureCloudTrailLogging automatic remediation action
- C. Create an AWS Config rule that is invoked when CloudTrail configuration changes Configure the rule to invoke an AWS Lambda function to enable CloudTrail
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) hourly rule with a schedule pattern to run an AWS Systems Manager Automation document to enable CloudTrail

**Answer: B**

#### NEW QUESTION 175

- (Exam Topic 1)

A company plans to launch a static website on its domain example.com and subdomain www.example.com using Amazon S3. How should the SysOps administrator meet this requirement?

- A. Create one S3 bucket named example.com for both the domain and subdomain.
- B. Create one S3 bucket with a wildcard named \*.example.com for both the domain and subdomain.
- C. Create two S3 buckets named example.com and www.example.com.
- D. Configure the subdomain bucket to redirect requests to the domain bucket.
- E. Create two S3 buckets named http://example.com and http://www.example.com.
- F. Configure the wildcard (\*) bucket to redirect requests to the domain bucket.

**Answer: C**

#### NEW QUESTION 176

- (Exam Topic 1)

A company runs its Infrastructure on Amazon EC2 Instances that run in an Auto Scaling group. Recently, the company promoted faulty code to the entire EC2 fleet. This faulty code caused the Auto Scaling group to scale the instances before any of the application logs could be retrieved. What should a SysOps administrator do to retain the application logs after instances are terminated?

- A. Configure an Auto Scaling lifecycle hook to create a snapshot of the ephemeral storage upon termination of the instances.
- B. Create a new Amazon Machine Image (AMI) that has the Amazon CloudWatch agent installed and configured to send logs to Amazon CloudWatch Log.
- C. Update the launch template to use the new AMI.
- D. Create a new Amazon Machine Image (AMI) that has a custom script configured to send logs to AWS CloudTrail.
- E. Update the launch template to use the new AMI.
- F. Install the Amazon CloudWatch agent on the Amazon Machine Image (AMI) that is defined in the launch template.
- G. Configure the CloudWatch agent to back up the logs to ephemeral storage.

**Answer: B**

#### NEW QUESTION 180

- (Exam Topic 1)

A company's SysOps administrator regularly checks the AWS Personal Health Dashboard in each of the company's accounts. The accounts are part of an organization in AWS Organizations. The company recently added 10 more accounts to the organization. The SysOps administrator must consolidate the alerts from each account's Personal Health Dashboard.

Which solution will meet this requirement with the LEAST amount of effort?

- A. Enable organizational view in AWS Health.
- B. Configure the Personal Health Dashboard in each account to forward events to a central AWS CloudTrail log.
- C. Create an AWS Lambda function to query the AWS Health API and to write all events to an Amazon DynamoDB table.
- D. Use the AWS Health API to write events to an Amazon DynamoDB table.

**Answer: A**

#### Explanation:

Enabling the organizational view in AWS Health will allow the SysOps administrator to consolidate the alerts from each account's Personal Health Dashboard. It will also provide the administrator with a single view of all the accounts in the organization, allowing them to easily monitor the health of all the accounts in the organization.

Reference:

[1] <https://aws.amazon.com/premiumsupport/knowledge-center/organizational-view-health-dashboard/>

#### NEW QUESTION 182

- (Exam Topic 1)

A SysOps administrator needs to configure a solution that will deliver digital content to a set of authorized users through Amazon CloudFront. Unauthorized users must be restricted from access. Which solution will meet these requirements?

- A. Store the digital content in an Amazon S3 bucket that does not have public access blocked.
- B. Use signed URLs to access the S3 bucket through CloudFront.
- C. Store the digital content in an Amazon S3 bucket that has public access blocked.
- D. Use an origin access identity (OAI) to deliver the content through CloudFront.
- E. Restrict S3 bucket access with signed URLs in CloudFront.
- F. Store the digital content in an Amazon S3 bucket that has public access blocked.
- G. Use an origin access identity (OAI) to deliver the content through CloudFront.
- H. Enable field-level encryption.
- I. Store the digital content in an Amazon S3 bucket that does not have public access blocked.
- J. Use signed cookies for restricted delivery of the content through CloudFront.

**Answer: B**

#### NEW QUESTION 187

- (Exam Topic 1)

A company has launched a social media website that gives users the ability to upload images directly to a centralized Amazon S3 bucket. The website is popular in areas that are geographically distant from the AWS Region where the S3 bucket is located. Users are reporting that uploads are slow. A SysOps administrator must improve the upload speed.

What should the SysOps administrator do to meet these requirements?

- A. Create S3 access points in Regions that are closer to the users.
- B. Create an accelerator in AWS Global Accelerator for the S3 bucket.
- C. Enable S3 Transfer Acceleration on the S3 bucket.
- D. Enable cross-origin resource sharing (CORS) on the S3 bucket.

**Answer:** C

**Explanation:**

You might want to use Transfer Acceleration on a bucket for various reasons: ->Your customers upload to a centralized bucket from all over the world. ->You transfer gigabytes to terabytes of data on a regular basis across continents. ->You can't use all of your available bandwidth over the internet when uploading to Amazon S3." <https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html>

**NEW QUESTION 189**

- (Exam Topic 1)

A SysOps administrator must configure a resilient tier of Amazon EC2 instances for a high performance computing (HPC) application. The HPC application requires minimum latency between nodes

Which actions should the SysOps administrator take to meet these requirements? (Select TWO.)

- A. Create an Amazon Elastic File System (Amazon EFS) file system Mount the file system to the EC2 instances by using user data
- B. Create a Multi-AZ Network Load Balancer in front of the EC2 instances
- C. Place the EC2 instances in an Auto Scaling group within a single subnet
- D. Launch the EC2 instances into a cluster placement group
- E. Launch the EC2 instances into a partition placement group

**Answer:** AD

**NEW QUESTION 190**

- (Exam Topic 1)

An application runs on multiple Amazon EC2 instances in an Auto Scaling group The Auto Scaling group is

configured to use the latest version of a launch template A SysOps administrator must devise a solution that centrally manages the application logs and retains the logs for no more than 90 days

Which solution will meet these requirements?

- A. Launch an Amazon Machine Image (AMI) that is preconfigured with the Amazon CloudWatch Logs agent to send logs to an Amazon S3 bucket Apply a 90-day S3 Lifecycle policy on the S3 bucket to expire the application logs
- B. Launch an Amazon Machine Image (AMI) that is preconfigured with the Amazon CloudWatch Logs agent to send logs to a log group Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled rule to perform an instance refresh every 90 days
- C. Update the launch template user data to install and configure the Amazon CloudWatch Logs agent to send logs to a log group Configure the retention period on the log group to be 90 days
- D. Update the launch template user data to install and configure the Amazon CloudWatch Logs agent to send logs to a log group Set the log rotation configuration of the EC2 instances to 90 days

**Answer:** C

**NEW QUESTION 193**

- (Exam Topic 1)

A company is using an Amazon DynamoDB table for data. A SysOps administrator must configure replication of the table to another AWS Region for disaster recovery.

What should the SysOps administrator do to meet this requirement?

- A. Enable DynamoDB Accelerator (DAX).
- B. Enable DynamoDB Streams, and add a global secondary index (GSI).
- C. Enable DynamoDB Streams, and-add a global table Region.
- D. Enable point-in-time recovery.

**Answer:** C

**NEW QUESTION 195**

- (Exam Topic 1)

A SysOps administrator needs to delete an AWS CloudFormation stack that is no longer in use. The CloudFormation stack is in the DELETE\_FAILED state. The SysOps administrator has validated the permissions that are required to delete the Cloud Formation stack.

- A. The configured timeout to delete the stack was too low for the delete operation to complete.
- B. The stack contains nested stacks that must be manually deleted fast.
- C. The stack was deployed with the -disable rollback option.
- D. There are additional resources associated with a security group in the stack
- E. There are Amazon S3 buckets that still contain objects in the stack.

**Answer:** DE

**NEW QUESTION 199**

- (Exam Topic 1)

A company hosts a web application on an Amazon EC2 instance. The web server logs are published to Amazon CloudWatch Logs. The log events have the same structure and include the HTTP response codes that are associated with the user requests. The company needs to monitor the number of times that the web server returns an HTTP 404 response.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create a CloudWatch Logs metric filter that counts the number of times that the web server returns an HTTP 404 response.
- B. Create a CloudWatch Logs subscription filter that counts the number of times that the web server returns an HTTP 404 response.
- C. Create an AWS Lambda function that runs a CloudWatch Logs Insights query that counts the number of 404 codes in the log events during the past hour.
- D. Create a script that runs a CloudWatch Logs Insights query that counts the number of 404 codes in the log events during the past hour.

**Answer:** A

**Explanation:**

This is the most operationally efficient solution that meets the requirements, as it will allow the company to monitor the number of times that the web server returns an HTTP 404 response in real-time. The other solutions (creating a CloudWatch Logs subscription filter, an AWS Lambda function, or a script) will require additional steps and resources to monitor the number of times that the web server returns an HTTP 404 response.

A metric filter allows you to search for specific terms, phrases, or values in your log events, and then to create a metric based on the number of occurrences of those search terms. This allows you to create a CloudWatch Metric that can be used to create alarms and dashboards, which can be used to monitor the number of HTTP 404 responses returned by the web server.

**NEW QUESTION 204**

- (Exam Topic 1)

A company has a critical serverless application that uses multiple AWS Lambda functions. Each Lambda function generates 1 GB of log data daily in its own Amazon CloudWatch Logs log group. The company's security team asks for a count of application errors, grouped by type, across all of the log groups.

What should a SysOps administrator do to meet this requirement?

- A. Perform a CloudWatch Logs Insights query that uses the stats command and count function.
- B. Perform a CloudWatch Logs search that uses the groupby keyword and count function.
- C. Perform an Amazon Athena query that uses the SELECT and GROUP BY keywords.
- D. Perform an Amazon RDS query that uses the SELECT and GROUP BY keywords.

**Answer:** A

**NEW QUESTION 209**

- (Exam Topic 1)

A Sysops administrator has created an Amazon EC2 instance using an AWS CloudFormation template in the us-east-1 Region. The administrator finds that this template has failed to create an EC2 instance in the us-west-2 Region. What is one cause for this failure?

- A. Resource tags defined in the CloudFormation template are specific to the us-east-1 Region.
- B. The Amazon Machine Image (AMI) ID referenced in the CloudFormation template could not be found in the us-west-2 Region.
- C. The cfn-init script did not run during resource provisioning in the us-west-2 Region.
- D. The IAM user was not created in the specified Region.

**Answer:** B

**Explanation:**

One possible cause for the failure of the CloudFormation template to create an EC2 instance in the us-west-2 Region is that the Amazon Machine Image (AMI) ID referenced in the template could not be found in the us-west-2 Region. This could be due to the fact that the AMI is not available in that region, or the credentials used to access the AMI were not configured properly. The other options (resource tags defined in the CloudFormation template are specific to the us-east-1 Region, the cfn-init script did not run during resource provisioning in the us-west-2 Region, and the IAM user was not created in the specified Region) are not valid causes for this failure.

**NEW QUESTION 213**

- (Exam Topic 1)

A company is managing multiple AWS accounts in AWS Organizations. The company is reviewing internal security of its AWS environment. The company's security administrator has their own AWS account and wants to review the VPC configuration of developer AWS accounts.

Which solution will meet these requirements in the MOST secure manner?

- A. Create an IAM policy in each developer account that has read-only access related to VPC resources Assign the policy to an IAM use
- B. Share the user credentials with the security administrator.
- C. Create an IAM policy in each developer account that has administrator access to all Amazon EC2 actions, including VPC action
- D. Assign the policy to an IAM user
- E. Share the user credentials with the security administrator.
- F. Create an IAM policy in each developer account that has administrator access related to VPC resources. Assign the policy to a cross-account IAM role
- G. Ask the security administrator to assume the role from their account.
- H. Create an IAM policy in each developer account that has read-only access related to VPC resources Assign the policy to a cross-account IAM role Ask the security administrator to assume the role from their account.

**Answer:** D

**NEW QUESTION 217**

- (Exam Topic 1)

A company's application currently uses an IAM role that allows all access to all AWS services. A SysOps administrator must ensure that the company's IAM policies allow only the permissions that the application requires.

How can the SysOps administrator create a policy to meet this requirement?

- A. Turn on AWS CloudTrail
- B. Generate a policy by using AWS Security Hub.
- C. Turn on Amazon EventBridge (Amazon CloudWatch Events). Generate a policy by using AWS Identity and Access Management Access Analyzer.
- D. Use the AWS CLI to run the get-generated-policy command in AWS Identity and Access Management Access Analyzer.
- E. Turn on AWS CloudTrail
- F. Generate a policy by using AWS Identity and Access Management Access Analyzer.

**Answer:** D

**Explanation:**

Generate a policy by using AWS Identity and Access Management Access Analyzer. AWS CloudTrail is a service that records all API calls made on your account. You can use this data to generate a policy with AWS Identity and Access Management Access Analyzer that only allows the permissions that the application requires. This will ensure that the application only has the necessary permissions and will protect the company from any unauthorized access.

<https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html#what-is-access-analyzer-poli>

#### NEW QUESTION 219

- (Exam Topic 1)

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.

**Answer: D**

#### NEW QUESTION 224

- (Exam Topic 1)

A company is using Amazon Elastic File System (Amazon EFS) to share a file system among several Amazon EC2 instances. As usage increases, users report that file retrieval from the EFS file system is slower than normal.

Which action should a SysOps administrator take to improve the performance of the file system?

- A. Configure the file system for Provisioned Throughput.
- B. Enable encryption in transit on the file system.
- C. Identify any unused files in the file system, and remove the unused files.
- D. Resize the Amazon Elastic Block Store (Amazon EBS) volume of each of the EC2 instances.

**Answer: A**

#### NEW QUESTION 225

- (Exam Topic 1)

A company asks a SysOps administrator to ensure that AWS CloudTrail files are not tampered with after they are created. Currently, the company uses AWS Identity and Access Management (IAM) to restrict access to specific trails. The company's security team needs the ability to trace the integrity of each file. What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function when a new file is delivered.
- B. Configure the Lambda function to compute an MD5 hash check on the file and store the result in an Amazon DynamoDB table.
- C. The security team can use the values that are stored in DynamoDB to verify the integrity of the delivered files.
- D. Create an AWS Lambda function that is invoked each time a new file is delivered to the CloudTrail bucket.
- E. Configure the Lambda function to compute an MD5 hash check on the file and store the result as a tag in an Amazon S3 object.
- F. The security team can use the information in the tag to verify the integrity of the delivered files.
- G. Enable the CloudTrail file integrity feature on an Amazon S3 bucket.
- H. Create an IAM policy that grants the security team access to the file integrity logs that are stored in the S3 bucket.
- I. Enable the CloudTrail file integrity feature on the trail.
- J. The security team can use the digest file that is created by CloudTrail to verify the integrity of the delivered files.

**Answer: D**

#### Explanation:

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html> "When you enable log file integrity validation, CloudTrail creates a hash for every log file that it delivers.

Every hour, CloudTrail also creates and delivers a file that references the log files for the last hour and contains a hash of each. This file is called a digest file. Validated log files are invaluable in security and forensic investigations"

#### NEW QUESTION 227

- (Exam Topic 1)

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 timeout set to 4 hours.
- C. DependsOn a timeout set to 4 hours.
- D. Metadata with a timeout set to 4 hours.

**Answer: B**

#### NEW QUESTION 229

- (Exam Topic 1)

A company runs several workloads on AWS. The company identifies five AWS Trusted Advisor service quota metrics to monitor in a specific AWS Region. The company wants to receive email notification each time resource usage exceeds 60% of one of the service quotas.

Which solution will meet these requirements?

- A. Create five Amazon CloudWatch alarms, one for each Trusted Advisor service quota metric.
- B. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification each time that usage exceeds 60% of one of the service quotas.
- C. Create five Amazon CloudWatch alarms, one for each Trusted Advisor service quota metric.
- D. Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification each time that usage exceeds 60% of one of the service quotas.
- E. Use the AWS Service Health Dashboard to monitor each Trusted Advisor service quota metric. Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification each time that usage exceeds 60% of one of the service quotas.
- F. Use the AWS Service Health Dashboard to monitor each Trusted Advisor service quota metric. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification each time that usage exceeds 60% of one of the service quotas.

topic for email notification each time that usage exceeds 60% of one of the service quotas.

**Answer:** A

**Explanation:**

CloudWatch alarms allow you to monitor AWS resources, and you can configure an SNS topic to send an email notification each time one of the alarms is triggered. This will ensure that the company receives email notifications each time one of the service quotas is exceeded, allowing the company to take action as needed.

**NEW QUESTION 231**

- (Exam Topic 1)

A company has multiple Amazon EC2 instances that run a resource-intensive application in a development environment. A SysOps administrator is implementing a solution to stop these EC2 instances when they are not in use.

Which solution will meet this requirement?

- A. Assess AWS CloudTrail logs to verify that there is no EC2 API activity
- B. Invoke an AWS Lambda function to stop the EC2 instances.
- C. Create an Amazon CloudWatch alarm to stop the EC2 instances when the average CPU utilization is lower than 5% for a 30-minute period.
- D. Create an Amazon CloudWatch metric to stop the EC2 instances when the VolumeReadBytes metric is lower than 500 for a 30-minute period.
- E. Use AWS Config to invoke an AWS Lambda function to stop the EC2 instances based on resource configuration changes.

**Answer:** B

**Explanation:**

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/UsingAlarmActions.html#AddingStopAction>

**NEW QUESTION 235**

- (Exam Topic 1)

A SysOps administrator is creating an Amazon EC2 Auto Scaling group in a new AWS account. After adding some instances, the SysOps administrator notices that the group has not reached the minimum number of instances. The SysOps administrator receives the following error message:

```
Launching a new EC2 instance. Status Reason: Your quota allows for 0 more running instance(s).  
You requested at least 1. Launching EC2 instance failed.
```

Which action will resolve this issue?

- A. Adjust the account spending limits for Amazon EC2 on the AWS Billing and Cost Management console
- B. Modify the EC2 quota for that AWS Region in the EC2 Settings section of the EC2 console.
- C. Request a quota Increase for the Instance type family by using Service Quotas on the AWS Management Console.
- D. Use the Rebalance action in the Auto Scaling group on the AWS Management Console.

**Answer:** C

**NEW QUESTION 236**

- (Exam Topic 1)

A SysOps administrator is unable to launch Amazon EC2 instances into a VPC because there are no available private IPv4 addresses in the VPC. Which combination of actions must the SysOps administrator take to launch the instances? (Select TWO.)

- A. Associate a secondary IPv4 CIDR block with the VPC
- B. Associate a primary IPv6 CIDR block with the VPC
- C. Create a new subnet for the VPC
- D. Modify the CIDR block of the VPC
- E. Modify the CIDR block of the subnet that is associated with the instances

**Answer:** AD

**NEW QUESTION 239**

- (Exam Topic 1)

A SysOps administrator is unable to authenticate an AWS CLI call to an AWS service. Which of the following is the cause of this issue?

- A. The IAM password is incorrect
- B. The server certificate is missing
- C. The SSH key pair is incorrect
- D. There is no access key

**Answer:** C

**NEW QUESTION 242**

- (Exam Topic 1)

A company uses an AWS CloudFormation template to provision an Amazon EC2 instance and an Amazon RDS DB instance. A SysOps administrator must update the template to ensure that the DB instance is created before the EC2 instance is launched.

What should the SysOps administrator do to meet this requirement?

- A. Add a wait condition to the template. Update the EC2 instance user data script to send a signal after the EC2 instance is started.
- B. Add the DependsOn attribute to the EC2 instance resource, and provide the logical name of the RDS resource.
- C. Change the order of the resources in the template so that the RDS resource is listed before the EC2 instance resource.
- D. Create multiple templates. Use AWS CloudFormation StackSets to wait for one stack to complete before the second stack is created.

**Answer:** B

**Explanation:**

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-dependson.html> Syntax The DependsOn attribute can take a single string or list of strings. "DependsOn" : [ String, ... ]

Example The following template contains an AWS::EC2::Instance resource with a DependsOn attribute that specifies myDB, an AWS::RDS::DBInstance. When CloudFormation creates this stack, it first creates myDB, then creates Ec2Instance.

**NEW QUESTION 245**

- (Exam Topic 1)

While setting up an AWS managed VPN connection, a SysOps administrator creates a customer gateway resource in AWS. The customer gateway device resides in a data center with a NAT gateway in front of it.

What address should be used to create the customer gateway resource?

- A. The private IP address of the customer gateway device
- B. The MAC address of the NAT device in front of the customer gateway device
- C. The public IP address of the customer gateway device
- D. The public IP address of the NAT device in front of the customer gateway device

**Answer:** D

**NEW QUESTION 248**

- (Exam Topic 1)

A company uses AWS CloudFormation to deploy its application infrastructure. Recently, a user accidentally changed a property of a database in a CloudFormation template and performed a stack update that caused an interruption to the application. A SysOps administrator must determine how to modify the deployment process to allow the DevOps team to continue to deploy the infrastructure, but prevent against accidental modifications to specific resources.

Which solution will meet these requirements?

- A. Set up an AWS Config rule to alert based on changes to any CloudFormation stack. An AWS Lambda function can then describe the stack to determine if any protected resources were modified and cancel the operation.
- B. Set up an Amazon CloudWatch Events event with a rule to trigger based on any CloudFormation API call. An AWS Lambda function can then describe the stack to determine if any protected resources were modified and cancel the operation.
- C. Launch the CloudFormation templates using a stack policy with an explicit allow for all resources and an explicit deny of the protected resources with an action of Update.
- D. Attach an IAM policy to the DevOps team role that prevents a CloudFormation stack from updating, with a condition based on the specific Amazon Resource Names (ARNs) of the protected resources.

**Answer:** B

**NEW QUESTION 250**

- (Exam Topic 1)

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

**Answer:** D

**Explanation:**

<https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>

**NEW QUESTION 255**

- (Exam Topic 1)

A SysOps administrator notices a scale-up event for an Amazon EC2 Auto Scaling group. Amazon CloudWatch shows a spike in the RequestCount metric for the associated Application Load Balancer. The administrator would like to know the IP addresses for the source of the requests.

Where can the administrator find this information?

- A. Auto Scaling logs
- B. AWS CloudTrail logs
- C. EC2 instance logs
- D. Elastic Load Balancer access logs

**Answer:** D

**Explanation:**

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and troubleshoot issues.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

**NEW QUESTION 258**

- (Exam Topic 1)

A company's web application is available through an Amazon CloudFront distribution and directly through an internet-facing Application Load Balancer (ALB). A

SysOps administrator must make the application accessible only through the CloudFront distribution and not directly through the ALB. The SysOps administrator must make this change without changing the application code  
Which solution will meet these requirements?

- A. Modify the ALB type to internal Set the distribution's origin to the internal ALB domain name
- B. Create a Lambda@Edge function Configure the function to compare a custom header value in the request with a stored password and to forward the request to the origin in case of a match Associate the function with the distribution.
- C. Replace the ALB with a new internal ALB Set the distribution's origin to the internal ALB domain name Add a custom HTTP header to the origin settings for the distribution In the ALB listener add a rule to forward requests that contain the matching custom header and the header's value Add a default rule to return a fixed response code of 403.
- D. Add a custom HTTP header to the origin settings for the distribution in the ALB listener add a rule to forward requests that contain the matching custom header and the header's value Add a default rule to return a fixed response code of 403.

**Answer: D**

**Explanation:**

To make the application accessible only through the CloudFront distribution and not directly through the Application Load Balancer (ALB), you can add a custom HTTP header to the origin settings for the CloudFront distribution. You can then create a rule in the ALB listener to forward requests that contain the matching custom header and its value to the origin. You can also add a default rule to the ALB listener to return a fixed response code of 403 for requests that do not contain the matching custom header. This will allow you to redirect all requests to the CloudFront distribution and block direct access to the application through the ALB.  
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>

**NEW QUESTION 262**

- (Exam Topic 1)

A SysOps administrator is setting up an automated process to recover an Amazon EC2 instance in the event of an underlying hardware failure. The recovered instance must have the same private IP address and the same Elastic IP address that the original instance had. The SysOps team must receive an email notification when the recovery process is initiated.

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm for the EC2 instance, and specify the StatusCheckFailedInstance metric
- B. Add an EC2 action to the alarm to recover the instance
- C. Add an alarm notification to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic
- D. Subscribe the SysOps team email address to the SNS topic.
- E. Create an Amazon CloudWatch alarm for the EC2 instance, and specify the StatusCheckFailed\_System metric
- F. Add an EC2 action to the alarm to recover the instance
- G. Add an alarm notification to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic
- H. Subscribe the SysOps team email address to the SNS topic.
- I. Create an Auto Scaling group across three different subnets in the same Availability Zone with a minimum, maximum, and desired size of 1. Configure the Auto Scaling group to use a launch template that specifies the private IP address and the Elastic IP address
- J. Add an activity notification for the Auto Scaling group to send an email message to the SysOps team through Amazon Simple Email Service (Amazon SES).
- K. Create an Auto Scaling group across three Availability Zones with a minimum, maximum, and desired size of 1. Configure the Auto Scaling group to use a launch template that specifies the private IP address and the Elastic IP address
- L. Add an activity notification for the Auto Scaling group to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic
- M. Subscribe the SysOps team email address to the SNS topic.

**Answer: B**

**Explanation:**

You can create an Amazon CloudWatch alarm that monitors an Amazon EC2 instance and automatically recovers the instance if it becomes impaired due to an underlying hardware failure or a problem that requires AWS involvement to repair. Terminated instances cannot be recovered. A recovered instance is identical to the original instance, including the instance ID, private IP addresses, Elastic IP addresses, and all instance metadata. If the impaired instance has a public IPv4 address, the instance retains the public IPv4 address after recovery. If the impaired instance is in a placement group, the recovered instance runs in the placement group. When the StatusCheckFailed\_System alarm is triggered, and the recover action is initiated, you will be notified by the Amazon SNS topic that you selected when you created the alarm and associated the recover action. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

**NEW QUESTION 265**

- (Exam Topic 1)

A company runs a web application on three Amazon EC2 instances behind an Application Load Balancer (ALB). The company notices that random periods of increased traffic cause a degradation in the application's performance. A SysOps administrator must scale the application to meet the increased traffic.

Which solution meets these requirements?

- A. Create an Amazon CloudWatch alarm to monitor application latency and increase the size of each EC2 instance if the desired threshold is reached.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor application latency and add an EC2 instance to the ALB if the desired threshold is reached.
- C. Deploy the application to an Auto Scaling group of EC2 instances with a target tracking scaling policy. Attach the ALB to the Auto Scaling group.
- D. Deploy the application to an Auto Scaling group of EC2 instances with a scheduled scaling policy. Attach the ALB to the Auto Scaling group.

**Answer: C**

**Explanation:**

[docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-target-tracking.html](https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-target-tracking.html)

**NEW QUESTION 266**

- (Exam Topic 1)

A SysOps administrator wants to upload a file that is 1 TB in size from on-premises to an Amazon S3 bucket using multipart uploads. What should the SysOps administrator do to meet this requirement?

- A. Upload the file using the S3 console.
- B. Use the s3api copy-object command.
- C. Use the s3api put-object command.

D. Use the s3 cp command.

**Answer:** D

**Explanation:**

It's a best practice to use aws s3 commands (such as aws s3 cp) for multipart uploads and downloads, because these aws s3 commands automatically perform multipart uploading and downloading based on the file size. By comparison, aws s3api commands, such as aws s3api create-multipart-upload, should be used only when aws s3 commands don't support a specific upload need, such as when the multipart upload involves multiple servers, a multipart upload is manually stopped and resumed later, or when the aws s3 command doesn't support a required request parameter.  
<https://aws.amazon.com/premiumsupport/knowledge-center/s3-multipart-upload-cli/>

**NEW QUESTION 270**

- (Exam Topic 1)

A SysOps administrator is troubleshooting connection timeouts to an Amazon EC2 instance that has a public IP address. The instance has a private IP address of 172.31.16.139. When the SysOps administrator tries to ping the instance's public IP address from the remote IP address 203.0.113.12, the response is "request timed out." The flow logs contain the following information:

```
2 123456789010 eni-1235b8ca123456789 203.0.113.12 172.31.16.139 0 0 1 4 336 1432917027 1432917142 ACCEPT OK
2 123456789010 eni-1235b8ca123456789 172.31.16.139 203.0.113.12 0 0 1 4 336 1432917094 1432917142 REJECT OK
```

What is one cause of the problem?

- A. Inbound security group deny rule
- B. Outbound security group deny rule
- C. Network ACL inbound rules
- D. Network ACL outbound rules

**Answer:** D

**NEW QUESTION 272**

- (Exam Topic 1)

A company is hosting applications on Amazon EC2 instances. The company is hosting a database on an Amazon RDS for PostgreSQL DB instance. The company requires all connections to the DB instance to be encrypted.

What should a SysOps administrator do to meet this requirement?

- A. Allow SSL connections to the database by using an inbound security group rule.
- B. Encrypt the database by using an AWS Key Management Service (AWS KMS) encryption key.
- C. Enforce SSL connections to the database by using a custom parameter group.
- D. Patch the database with SSL/TLS by using a custom PostgreSQL extension.

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/PostgreSQL.Concepts.General.SSL.htm> Amazon RDS supports SSL/TLS encryption for connections to the database, and this can be enabled by creating a custom parameter group and setting the rds.force\_ssl parameter to 1. This will ensure that all connections to the database are encrypted, protecting the data and maintaining compliance with the company's requirements.

**NEW QUESTION 276**

- (Exam Topic 1)

A company with multiple AWS accounts needs to obtain recommendations for AWS Lambda functions and identify optimal resource configurations for each Lambda function. How should a SysOps administrator provide these recommendations?

- A. Create an AWS Serverless Application Repository and export the Lambda function recommendations.
- B. Enable AWS Compute Optimizer and export the Lambda function recommendations
- C. Enable all features of AWS Organization and export the recommendations from AWS CloudTrailInsights.
- D. Run AWS Trusted Advisor and export the Lambda function recommendations

**Answer:** B

**NEW QUESTION 280**

- (Exam Topic 1)

A company is running an application on premises and wants to use AWS for data backup All of the data must be available locally The backup application can write only to block-based storage that is compatible with the Portable Operating System Interface (POSIX)

Which backup solution will meet these requirements?

- A. Configure the backup software to use Amazon S3 as the target for the data backups
- B. Configure the backup software to use Amazon S3 Glacier as the target for the data backups
- C. Use AWS Storage Gateway, and configure it to use gateway-cached volumes
- D. Use AWS Storage Gateway, and configure it to use gateway-stored volumes

**Answer:** D

**Explanation:**

<https://docs.aws.amazon.com/storagegateway/latest/userguide/StorageGatewayConcepts.html>

#### NEW QUESTION 285

- (Exam Topic 1)

A company has an internal web application that runs on Amazon EC2 instances behind an Application Load

Balancer. The instances run in an Amazon EC2 Auto Scaling group in a single Availability Zone. A SysOps administrator must make the application highly available.

Which action should the SysOps administrator take to meet this requirement?

- A. Increase the maximum number of instances in the Auto Scaling group to meet the capacity that is required at peak usage.
- B. Increase the minimum number of instances in the Auto Scaling group to meet the capacity that is required at peak usage.
- C. Update the Auto Scaling group to launch new instances in a second Availability Zone in the same AWS Region.
- D. Update the Auto Scaling group to launch new instances in an Availability Zone in a second AWS Region.

**Answer: C**

#### NEW QUESTION 290

- (Exam Topic 1)

A company uses an Amazon S3 bucket to store data files. The S3 bucket contains hundreds of objects. The company needs to replace a tag on all the objects in the S3 bucket with another tag.

What is the MOST operationally efficient way to meet this requirement?

- A. Use S3 Batch Operation
- B. Specify the operation to replace all object tags.
- C. Use the AWS CLI to get the tags for each object
- D. Save the tags in a list
- E. Use S3 Batch Operations. Specify the operation to delete all object tags
- F. Use the AWS CLI and the list to retag the objects.
- G. Use the AWS CLI to get the tags for each object
- H. Save the tags in a list
- I. Use the AWS CLI and the list to remove the object tags
- J. Use the AWS CLI and the list to retag the objects.
- K. Use the AWS CLI to copy the objects to another S3 bucket
- L. Add the new tag to the copied objects. Delete the original objects.

**Answer: A**

#### Explanation:

Ref. <https://aws.amazon.com/es/blogs/storage/adding-and-removing-object-tags-with-s3-batch-operations/>

#### NEW QUESTION 292

- (Exam Topic 1)

A company uses AWS Organizations to manage its AWS accounts. A SysOps administrator must create a backup strategy for all Amazon EC2 instances across all the company's AWS accounts.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Deploy an AWS Lambda function to each account to run EC2 instance snapshots on a scheduled basis.
- B. Create an AWS CloudFormation stack set in the management account to add an AutoBackup=True tag to every EC2 instance
- C. Use AWS Backup in the management account to deploy policies for all accounts and resources.
- D. Use a service control policy (SCP) to run EC2 instance snapshots on a scheduled basis in each account.

**Answer: B**

#### NEW QUESTION 293

- (Exam Topic 1)

A company recently migrated its server infrastructure to Amazon EC2 instances. The company wants to use Amazon CloudWatch metrics to track instance memory utilization and available disk space.

What should a SysOps administrator do to meet these requirements?

- A. Configure CloudWatch from the AWS Management Console for all the instances that require monitoring by CloudWatch
- B. AWS automatically installs and configures the agents for the specified instances.
- C. Install and configure the CloudWatch agent on all the instances
- D. Attach an IAM role to allow the instances to write logs to CloudWatch.
- E. Install and configure the CloudWatch agent on all the instances
- F. Attach an IAM user to allow the instances to write logs to CloudWatch.
- G. Install and configure the CloudWatch agent on all the instances
- H. Attach the necessary security groups to allow the instances to write logs to CloudWatch

**Answer: C**

#### NEW QUESTION 295

- (Exam Topic 1)

A company needs to implement a managed file system to host Windows file shares for users on premises. Resources in the AWS Cloud also need access to the data on these file shares. A SysOps administrator needs to present the user file shares on premises and make the user file shares available on AWS with minimum latency.

What should the SysOps administrator do to meet these requirements?

- A. Set up an Amazon S3 File Gateway.
- B. Set up an AWS Direct Connect connection.
- C. Use AWS DataSync to automate data transfers between the existing file servers and AWS.
- D. Set up an Amazon FSx File Gateway.

**Answer:** D

**Explanation:**

Amazon FSx provides a fully managed file system that is optimized for Windows-based workloads and can be used to create file shares that can be accessed both on premises and in the AWS Cloud. The file shares that are created in Amazon FSx are highly available and can be accessed with low latency. Additionally, Amazon FSx supports Windows-based authentication, making it easy to integrate with existing Windows user accounts.

References:

[1] <https://aws.amazon.com/fsx/>

[2] <https://aws.amazon.com/storage/file-storage/>

[3] <https://docs.aws.a>

**NEW QUESTION 297**

- (Exam Topic 1)

A SysOps administrator needs to create alerts that are based on the read and write metrics of Amazon Elastic Block Store (Amazon EBS) volumes that are attached to an Amazon EC2 instance. The SysOps administrator creates and enables Amazon CloudWatch alarms for the DiskReadBytes metric and the DiskWriteBytes metric.

A custom monitoring tool that is installed on the EC2 instance with the same alarm configuration indicates that the volume metrics have exceeded the threshold. However, the CloudWatch alarms were not in ALARM state.

Which action will ensure that the CloudWatch alarms function correctly?

- A. Install and configure the CloudWatch agent on the EC2 instance to capture the desired metrics.
- B. Install and configure AWS Systems Manager Agent on the EC2 instance to capture the desired metrics.
- C. Reconfigure the CloudWatch alarms to use the VolumeReadBytes metric and the VolumeWriteBytes metric for the EBS volumes.
- D. Reconfigure the CloudWatch alarms to use the VolumeReadBytes metric and the VolumeWriteBytes metric for the EC2 instance.

**Answer:** A

**NEW QUESTION 299**

- (Exam Topic 1)

A company has an existing web application that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB) across two Availability Zones. The application uses an Amazon RDS Multi-AZ DB Instance. Amazon Route 53 record sets route requests for dynamic content to the load balancer and requests for static content to an Amazon S3 bucket. Site visitors are reporting extremely long loading times.

Which actions should be taken to improve the performance of the website? (Select TWO )

- A. Add Amazon CloudFront caching for static content
- B. Change the load balancer listener from HTTPS to TCP
- C. Enable Amazon Route 53 latency-based routing
- D. Implement Amazon EC2 Auto Scaling for the web servers
- E. Move the static content from Amazon S3 to the web servers

**Answer:** AD

**NEW QUESTION 301**

- (Exam Topic 1)

A company runs a stateless application that is hosted on an Amazon EC2 instance. Users are reporting performance issues. A SysOps administrator reviews the Amazon CloudWatch metrics for the application and notices that the instance's CPU utilization frequently reaches 90% during business hours.

What is the MOST operationally efficient solution that will improve the application's responsiveness?

- A. Configure CloudWatch logging on the EC2 instance
- B. Configure a CloudWatch alarm for CPU utilization to alert the SysOps administrator when CPU utilization goes above 90%.
- C. Configure an AWS Client VPN connection to allow the application users to connect directly to the EC2 instance private IP address to reduce latency.
- D. Create an Auto Scaling group, and assign it to an Application Load Balance
- E. Configure a target tracking scaling policy that is based on the average CPU utilization of the Auto Scaling group.
- F. Create a CloudWatch alarm that activates when the EC2 instance's CPU utilization goes above 80%. Configure the alarm to invoke an AWS Lambda function that vertically scales the instance.

**Answer:** C

**NEW QUESTION 306**

- (Exam Topic 1)

A company is partnering with an external vendor to provide data processing services. For this integration, the vendor must host the company's data in an Amazon S3 bucket in the vendor's AWS account. The vendor is allowing the company to provide an AWS Key Management Service (AWS KMS) key to encrypt the company's data. The vendor has provided an IAM role Amazon Resource Name (ARN) to the company for this integration.

What should a SysOps administrator do to configure this integration?

- A. Create a new KMS key
- B. Add the vendor's IAM role ARN to the KMS key policy
- C. Provide the new KMS key ARN to the vendor.
- D. Create a new KMS key
- E. Create a new IAM user
- F. Add the vendor's IAM role ARN to an inline policy that is attached to the IAM user
- G. Provide the new IAM user ARN to the vendor.
- H. Configure encryption using the KMS managed S3 key
- I. Add the vendor's IAM role ARN to the KMS managed S3 key policy
- J. Provide the KMS managed S3 key ARN to the vendor.
- K. Configure encryption using the KMS managed S3 key
- L. Create an S3 bucket
- M. Add the vendor's IAM role ARN to the S3 bucket policy
- N. Provide the S3 bucket ARN to the vendor.

**Answer: C**

**NEW QUESTION 308**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

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

2nd - Questions and Answers in PDF Format

### **AWS-SysOps Practice Exam Features:**

- \* AWS-SysOps Questions and Answers Updated Frequently
- \* AWS-SysOps Practice Questions Verified by Expert Senior Certified Staff
- \* AWS-SysOps Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* AWS-SysOps Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

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