

Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional

<https://www.2passeasy.com/dumps/DOP-C02/>



NEW QUESTION 1

A company's DevOps engineer uses AWS Systems Manager to perform maintenance tasks during maintenance windows. The company has a few Amazon EC2 instances that require a restart after notifications from AWS Health. The DevOps engineer needs to implement an automated solution to remediate these notifications. The DevOps engineer creates an Amazon EventBridge rule.

How should the DevOps engineer configure the EventBridge rule to meet these requirements?

- A. Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance.
- B. Target a Systems Manager document to restart the EC2 instance.
- C. Configure an event source of Systems Manager and an event type that indicates a maintenance window.
- D. Target a Systems Manager document to restart the EC2 instance.
- E. Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance.
- F. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.
- G. Configure an event source of EC2 and an event type that indicates instance maintenance.
- H. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.

Answer: C

Explanation:

AWS Health provides real-time events and information related to your AWS infrastructure. It can be integrated with Amazon EventBridge to act upon the health events automatically. If the maintenance notification from AWS Health indicates that an EC2 instance requires a restart, you can set up an EventBridge rule to respond to such events. In this case, the target of this rule would be a Lambda function that would trigger a Systems Manager automation to restart the EC2 instance during a maintenance window. Remember, AWS Health is the source of the events (not EC2 or Systems Manager), and AWS Lambda can be used to execute complex remediation tasks, such as scheduling maintenance tasks via Systems Manager.

The following are the steps involved in configuring the EventBridge rule to meet these requirements:

? Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance.

? Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.

The AWS Lambda function will be triggered by the event from AWS Health. The function will then register an automation task to restart the EC2 instance during the next maintenance window.

NEW QUESTION 2

A company runs applications in AWS accounts that are in an organization in AWS Organizations. The applications use Amazon EC2 instances and Amazon S3. The company wants to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future. When the company detects one of these events, the company wants to use an existing Amazon Simple Notification Service (Amazon SNS) topic to send a notification to its operational support team for investigation and remediation.

Which solution will meet these requirements in accordance with AWS best practices?

- A. In the organization's management account, configure an AWS account as the Amazon GuardDuty administrator account.
- B. In the GuardDuty administrator account, add the company's existing AWS accounts to GuardDuty as members. In the GuardDuty administrator account, create an Amazon EventBridge rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic.
- C. In the organization's management account, configure Amazon GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts. Create an AWS CloudFormation stack set that accepts the GuardDuty invitation and creates an Amazon EventBridge rule. Configure the rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic.
- D. GuardDuty events and to forward matching events to the SNS topic.
- E. Configure the CloudFormation stack set to deploy into all AWS accounts in the organization.
- F. In the organization's management account, create an AWS CloudTrail organization trail. Activate the organization trail in all AWS accounts in the organization.
- G. Create an SCP that enables VPC Flow Logs in each account in the organization.
- H. Create an SCP that enables VPC Flow Logs in each account in the organization.
- I. Configure AWS Security Hub for the organization. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.
- J. In the organization's management account, configure an AWS account as the AWS CloudTrail administrator account. In the CloudTrail administrator account, create a CloudTrail organization trail.
- K. Add the company's existing AWS accounts to the organization trail. Create an SCP that enables VPC Flow Logs in each account in the organization.
- L. Configure AWS Security Hub for the organization.
- M. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.

Answer: B

Explanation:

It allows the company to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future using Amazon GuardDuty. It also provides a solution for automatically adding future AWS accounts to GuardDuty by configuring GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts.

NEW QUESTION 3

A company uses a single AWS account to test applications on Amazon EC2 instances. The company has turned on AWS Config in the AWS account and has activated the restricted-ssh AWS Config managed rule.

The company needs an automated monitoring solution that will provide a customized notification in real time if any security group in the account is not compliant with the restricted-ssh rule. The customized notification must contain the name and ID of the noncompliant security group.

A DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic in the account and subscribes the appropriate personnel to the topic.

What should the DevOps engineer do next to meet these requirements?

- A. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule.
- B. Configure an input transformer for the EventBridge rule. Configure the EventBridge rule to publish a notification to the SNS topic.
- C. Configure AWS Config to send all evaluation results for the restricted-ssh rule to the SNS topic.
- D. Configure a filter policy on the SNS topic to send only notifications that contain the text of NON_COMPLIANT in the notification to subscribers.
- E. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure the EventBridge rule to invoke AWS Systems Manager Run Command on the SNS topic to customize a notification and to publish the notification to the SNS topic.
- F. Create an Amazon EventBridge rule that matches all AWS Config evaluation results of NON_COMPLIANT. Configure an input transformer for the restricted-ssh rule. Configure the EventBridge rule to publish a notification to the SNS topic.

Answer: A

Explanation:

Create an Amazon EventBridge (Amazon CloudWatch Events) rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure an input transformer for the EventBridge (CloudWatch Events) rule. Configure the EventBridge (CloudWatch Events) rule to publish a notification to the SNS topic. This approach uses Amazon EventBridge (previously known as Amazon CloudWatch Events) to filter AWS Config evaluation results based on the restricted-ssh rule and its compliance status (NON_COMPLIANT). An input transformer can be used to customize the information contained in the notification, such as the name and ID of the noncompliant security group. The EventBridge (CloudWatch Events) rule can then be configured to publish a notification to the SNS topic, which will notify the appropriate personnel in real-time.

NEW QUESTION 4

A DevOps engineer is designing an application that integrates with a legacy REST API. The application has an AWS Lambda function that reads records from an Amazon Kinesis data stream. The Lambda function sends the records to the legacy REST API.

Approximately 10% of the records that the Lambda function sends from the Kinesis data stream have data errors and must be processed manually. The Lambda function event source configuration has an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as an on-failure destination. The DevOps engineer has configured the Lambda function to process records in batches and has implemented retries in case of failure.

During testing the DevOps engineer notices that the dead-letter queue contains many records that have no data errors and that already have been processed by the legacy REST API. The DevOps engineer needs to configure the Lambda function's event source options to reduce the number of errorless records that are sent to the dead-letter queue.

Which solution will meet these requirements?

- A. Increase the retry attempts
- B. Configure the setting to split the batch when an error occurs
- C. Increase the concurrent batches per shard
- D. Decrease the maximum age of record

Answer: B

Explanation:

This solution will meet the requirements because it will reduce the number of errorless records that are sent to the dead-letter queue. When you configure the setting to split the batch when an error occurs, Lambda will retry only the records that caused the error, instead of retrying the entire batch. This way, the records that have no data errors and have already been processed by the legacy REST API will not be retried and sent to the dead-letter queue unnecessarily.

<https://docs.aws.amazon.com/lambda/latest/dg/with-kinesis.html>

NEW QUESTION 5

A company runs an application on one Amazon EC2 instance. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance must restart or relaunch automatically if the instance becomes unresponsive.

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm for the StatusCheckFailed metri
- B. Use the recover action to stop and start the instanc
- C. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
- D. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instanc
- E. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.
- F. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failur
- G. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
- H. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resourc
- I. Add a command in UserData to retrieve the application metadata from Amazon S3.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/mt/how-to-set-up-aws-opsworks-stacks-auto-healing-notifications-in-amazon-cloudwatch-events/>

NEW QUESTION 6

A company deploys a web application on Amazon EC2 instances that are behind an

Application Load Balancer (ALB). The company stores the application code in an AWS CodeCommit repository. When code is merged to the main branch, an AWS Lambda function invokes an AWS CodeBuild project. The CodeBuild project packages the code, stores the packaged code in AWS CodeArtifact, and invokes AWS Systems Manager Run Command to deploy the packaged code to the EC2 instances.

Previous deployments have resulted in defects, EC2 instances that are not running the latest version of the packaged code, and inconsistencies between instances.

Which combination of actions should a DevOps engineer take to implement a more reliable deployment solution? (Select TWO.)

- A. Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provide
- B. Configure pipeline stages that run the CodeBuild project in parallel to build and test the applicatio
- C. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.
- D. Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provide
- E. Create separate pipeline stages that run a CodeBuild project to build and then test the applicatio
- F. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.
- G. Create an AWS CodeDeploy application and a deployment group to deploy the packaged code to the EC2 instance
- H. Configure the ALB for the deployment group.
- I. Create individual Lambda functions that use AWS CodeDeploy instead of Systems Manager to run build, test, and deploy actions.
- J. Create an Amazon S3 bucke
- K. Modify the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifac
- L. Use deploy actions in CodeDeploy to deploy the artifact to the EC2 instances.

Answer: AC

Explanation:

To implement a more reliable deployment solution, a DevOps engineer should take the following actions:

? Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provider. Configure pipeline stages that run the CodeBuild project in parallel to build and test the application. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action. This action will improve the deployment reliability by automating the entire process from code commit to deployment, reducing human errors and inconsistencies. By running the build and test stages in parallel, the pipeline can also speed up the delivery time and provide faster feedback. By using CodeDeploy as the deployment action, the pipeline can leverage the features of CodeDeploy, such as traffic shifting, health checks, rollback, and deployment configuration¹²³

? Create an AWS CodeDeploy application and a deployment group to deploy the packaged code to the EC2 instances. Configure the ALB for the deployment group. This action will improve the deployment reliability by using CodeDeploy to orchestrate the deployment across multiple EC2 instances behind an ALB. CodeDeploy can perform blue/green deployments or in-place deployments with traffic shifting, which can minimize downtime and reduce risks. CodeDeploy can also monitor the health of the instances during and after the deployment, and automatically roll back if any issues are detected. By configuring the ALB for the deployment group, CodeDeploy can register and deregister instances from the load balancer as needed, ensuring that only healthy instances receive traffic⁴⁵

The other options are not correct because they do not improve the deployment reliability or follow best practices. Creating separate pipeline stages that run a CodeBuild project to build and then test the application is not a good option because it will increase the pipeline execution time and delay the feedback loop. Creating individual Lambda functions that use CodeDeploy instead of Systems Manager to run build, test, and deploy actions is not a valid option because it will add unnecessary complexity and cost to the solution. Lambda functions are not designed for long-running tasks such as building or deploying applications. Creating an Amazon S3 bucket and modifying the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifact is not a necessary option because it will not affect the deployment reliability. CodeArtifact is a secure, scalable, and cost-effective package management service that can store and share software packages for application development⁶⁷

References:

- ? 1: What is AWS CodePipeline? - AWS CodePipeline
- ? 2: Create a pipeline in AWS CodePipeline - AWS CodePipeline
- ? 3: Deploy an application with AWS CodeDeploy - AWS CodePipeline
- ? 4: What is AWS CodeDeploy? - AWS CodeDeploy
- ? 5: Configure an Application Load Balancer for your blue/green deployments - AWS CodeDeploy
- ? 6: What is AWS Lambda? - AWS Lambda
- ? 7: What is AWS CodeArtifact? - AWS CodeArtifact

NEW QUESTION 7

A DevOps engineer is building an application that uses an AWS Lambda function to query an Amazon Aurora MySQL DB cluster. The Lambda function performs only read queries. Amazon EventBridge events invoke the Lambda function.

As more events invoke the Lambda function each second, the database's latency increases and the database's throughput decreases. The DevOps engineer needs to improve the performance of the application.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Use Amazon RDS Proxy to create a proxy
- B. Connect the proxy to the Aurora cluster reader endpoint
- C. Set a maximum connections percentage on the proxy.
- D. Implement database connection pooling inside the Lambda code
- E. Set a maximum number of connections on the database connection pool.
- F. Implement the database connection opening outside the Lambda event handler code.
- G. Implement the database connection opening and closing inside the Lambda event handler code.
- H. Connect to the proxy endpoint from the Lambda function.
- I. Connect to the Aurora cluster endpoint from the Lambda function.

Answer: ACE

Explanation:

To improve the performance of the application, the DevOps engineer should use Amazon RDS Proxy, implement the database connection opening outside the Lambda event handler code, and connect to the proxy endpoint from the Lambda function. References:

? Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure¹. By using Amazon RDS Proxy, the DevOps engineer can reduce the overhead of opening and closing connections to the database, which can improve latency and throughput².

? The DevOps engineer should connect the proxy to the Aurora cluster reader endpoint, which allows read-only connections to one of the Aurora Replicas in the DB cluster³. This can help balance the load across multiple read replicas and improve performance for read-intensive workloads⁴.

? The DevOps engineer should implement the database connection opening outside the Lambda event handler code, which means using a global variable to store the database connection object⁵. This can enable connection reuse across multiple invocations of the Lambda function, which can reduce latency and improve performance.

? The DevOps engineer should connect to the proxy endpoint from the Lambda function, which is a unique URL that represents the proxy. This can allow the Lambda function to access the database through the proxy, which can provide benefits such as connection pooling, load balancing, failover handling, and enhanced security.

? The other options are incorrect because:

NEW QUESTION 8

A company is implementing a well-architected design for its globally accessible API stack. The design needs to ensure both high reliability and fast response times for users located in North America and Europe.

The API stack contains the following three tiers: Amazon API Gateway

AWS Lambda Amazon DynamoDB

Which solution will meet the requirements?

- A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using health check
- B. Configure the APIs to forward requests to a Lambda function in that Region
- C. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.
- D. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using latency-based routing and health check
- E. Configure the APIs to forward requests to a Lambda function in that Region
- F. Configure the Lambda functions to retrieve and update the data in a DynamoDB global table.
- G. Configure Amazon Route 53 to point to API Gateway in North America, create a disaster recovery API in Europe, and configure both APIs to forward requests to the Lambda functions in that Region
- H. Retrieve the data from a DynamoDB global table
- I. Deploy a Lambda function to check the North America API health every 5 minutes
- J. In the event of a failure, update Route 53 to point to the disaster recovery API.
- K. Configure Amazon Route 53 to point to API Gateway API in North America using latency-based routing

- L. Configure the API to forward requests to the Lambda function in the Region nearest to the use
- M. Configure the Lambda function to retrieve and update the data in a DynamoDB table.

Answer: B

NEW QUESTION 9

A company has a mobile application that makes HTTP API calls to an Application Load Balancer (ALB). The ALB routes requests to an AWS Lambda function. Many different versions of the application are in use at any given time, including versions that are in testing by a subset of users. The version of the application is defined in the user-agent header that is sent with all requests to the API.

After a series of recent changes to the API, the company has observed issues with the application. The company needs to gather a metric for each API operation by response code for each version of the application that is in use. A DevOps engineer has modified the Lambda function to extract the API operation name, version information from the user-agent header and response code.

Which additional set of actions should the DevOps engineer take to gather the required metrics?

- A. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group
- B. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name
- C. Specify response code and application version as dimensions for the metric.
- D. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group
- E. Configure a CloudWatch Logs Insights query to populate CloudWatch metrics from the log line
- F. Specify response code and application version as dimensions for the metric.
- G. Configure the ALB access logs to write to an Amazon CloudWatch Logs log group
- H. Modify the Lambda function to respond to the ALB with the API operation name, response code, and version number as response metadata
- I. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name
- J. Specify response code and application version as dimensions for the metric.
- K. Configure AWS X-Ray integration on the Lambda function
- L. Modify the Lambda function to create an X-Ray subsegment with the API operation name, response code, and version number
- M. Configure X-Ray insights to extract an aggregated metric for each API operation name and to publish the metric to Amazon CloudWatch
- N. Specify response code and application version as dimensions for the metric.

Answer: A

Explanation:

"Note that the metric filter is different from a log insights query, where the experience is interactive and provides immediate search results for the user to investigate.

No automatic action can be invoked from an insights query. Metric filters, on the other hand, will generate metric data in the form of a time series. This lets you create alarms that integrate into your ITSM processes, execute AWS Lambda functions, or even create anomaly detection models."

<https://aws.amazon.com/blogs/mt/quantify-custom-application-metrics-with-amazon-cloudwatch-logs-and-metric-filters/>

NEW QUESTION 10

A growing company manages more than 50 accounts in an organization in AWS Organizations. The company has configured its applications to send logs to Amazon CloudWatch Logs.

A DevOps engineer needs to aggregate logs so that the company can quickly search the logs to respond to future security incidents. The DevOps engineer has created a new AWS account for centralized monitoring.

Which combination of steps should the DevOps engineer take to make the application logs searchable from the monitoring account? (Select THREE.)

- A. In the monitoring account, download an AWS CloudFormation template from CloudWatch to use in Organization
- B. Use CloudFormation StackSets in the organization's management account to deploy the CloudFormation template to the entire organization.
- C. Create an AWS CloudFormation template that defines an IAM role
- D. Configure the role to allow logs.amazonaws.com to perform the logs:Link action if the aws:ResourceAccount property is equal to the monitoring account ID
- E. Use CloudFormation StackSets in the organization's management account to deploy the CloudFormation template to the entire organization.
- F. Create an IAM role in the monitoring account
- G. Attach a trust policy that allows logs.amazonaws.com to perform the iam:CreateSink action if the aws:PrincipalOrgId property is equal to the organization ID.
- H. In the organization's management account, enable the logging policies for the organization.
- I. Use CloudWatch Observability Access Manager in the monitoring account to create a sink
- J. Allow logs to be shared with the monitoring account
- K. Configure the monitoring account data selection to view the Observability data from the organization ID.
- L. In the monitoring account, attach the CloudWatchLogsReadOnlyAccess AWS managed policy to an IAM role that can be assumed to search the logs.

Answer: BCF

Explanation:

? To aggregate logs from multiple accounts in an organization, the DevOps engineer needs to create a cross-account subscription¹ that allows the monitoring account to receive log events from the sharing accounts.

? To enable cross-account subscription, the DevOps engineer needs to create an IAM role in each sharing account that grants permission to CloudWatch Logs to link the log groups to the destination in the monitoring account². This can be done using a CloudFormation template and StackSets³ to deploy the role to all accounts in the organization.

? The DevOps engineer also needs to create an IAM role in the monitoring account that allows CloudWatch Logs to create a sink for receiving log events from other accounts⁴. The role must have a trust policy that specifies the organization ID as a condition.

? Finally, the DevOps engineer needs to attach the CloudWatchLogsReadOnlyAccess policy⁵ to an IAM role in the monitoring account that can be used to search the logs from the cross-account subscription.

References: 1: Cross-account log data sharing with subscriptions 2: Create an IAM role for CloudWatch Logs in each sharing account 3: AWS CloudFormation StackSets 4: Create an IAM role for CloudWatch Logs in your monitoring account 5: CloudWatchLogsReadOnlyAccess policy

NEW QUESTION 10

A DevOps engineer has automated a web service deployment by using AWS CodePipeline with the following steps:

- 1) An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
- 2) An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
- 3) A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment.

The quality assurance (QA) team requests permission to inspect the build artifact before the deployment to the production environment occurs. The QA team wants to run an internal penetration testing tool to conduct manual tests. The tool will be invoked by a REST API call.

Which combination of actions should the DevOps engineer take to fulfill this request? (Choose two.)

- A. Insert a manual approval action between the test actions and deployment actions of the pipeline.
- B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.
- C. Update the CodeDeploy deployment groups so that they require manual approval to proceed.
- D. Update the pipeline to directly call the REST API for the penetration testing tool.
- E. Update the pipeline to invoke an AWS Lambda function that calls the REST API for the penetration testing tool.

Answer: AE

NEW QUESTION 12

A company uses AWS CodePipeline pipelines to automate releases of its application. A typical pipeline consists of three stages: build, test, and deployment. The company has been using a separate AWS CodeBuild project to run scripts for each stage. However, the company now wants to use AWS CodeDeploy to handle the deployment stage of the pipelines.

The company has packaged the application as an RPM package and must deploy the application to a fleet of Amazon EC2 instances. The EC2 instances are in an EC2 Auto Scaling group and are launched from a common AMI.

Which combination of steps should a DevOps engineer perform to meet these requirements? (Choose two.)

- A. Create a new version of the common AMI with the CodeDeploy agent installed
- B. Update the IAM role of the EC2 instances to allow access to CodeDeploy.
- C. Create a new version of the common AMI with the CodeDeploy agent installed
- D. Create an AppSpec file that contains application deployment scripts and grants access to CodeDeploy.
- E. Create an application in CodeDeploy
- F. Configure an in-place deployment type
- G. Specify the Auto Scaling group as the deployment target
- H. Add a step to the CodePipeline pipeline to use EC2 Image Builder to create a new AMI
- I. Configure CodeDeploy to deploy the newly created AMI.
- J. Create an application in CodeDeploy
- K. Configure an in-place deployment type
- L. Specify the Auto Scaling group as the deployment target
- M. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.
- N. Create an application in CodeDeploy
- O. Configure an in-place deployment type
- P. Specify the EC2 instances that are launched from the common AMI as the deployment target
- Q. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.

Answer: AD

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html>

NEW QUESTION 13

A company recently launched multiple applications that use Application Load Balancers. Application response time often slows down when the applications experience problems. A DevOps engineer needs to implement a monitoring solution that alerts the company when the applications begin to perform slowly. The DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic and subscribes the company's email address to the topic. What should the DevOps engineer do next to meet the requirements?

- A. Create an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval. Configure the Lambda function to publish a notification to the SNS topic when the applications return errors.
- B. Create an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interval.
- C. Configure the canary to use the SNS topic when the applications return errors.
- D. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric. Configure the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports. Configure the CloudWatch alarm to use the SNS topic.
- E. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric. Configure the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports. Configure the CloudWatch alarm to use the SNS topic.

Answer: B

Explanation:

? Option A is incorrect because creating an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval is not a valid solution. EventBridge rules can only trigger Lambda functions based on events, not on time intervals. Moreover, querying the applications on a 5-minute interval might incur unnecessary costs and network overhead, and might not detect performance issues in real time.

? Option B is correct because creating an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interval is a valid solution. CloudWatch Synthetics canaries are configurable scripts that monitor endpoints and APIs by simulating customer behavior. Canaries can run as often as once per minute, and can measure the latency and availability of the applications. Canaries can also send notifications to an Amazon SNS topic when they detect errors or performance issues¹.

? Option C is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution. The RequestCountPerTarget metric measures the number of requests completed or connections made per target in a target group². This metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports is not a valid way to measure the application performance, as it depends on the application design and implementation.

? Option D is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution, for the same reason as option C. The RequestCountPerTarget metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports is not a valid way to measure the application performance, as it does not account for variability or outliers in the response time distribution.

References:

? 1: Using synthetic monitoring

? 2: Application Load Balancer metrics

NEW QUESTION 14

A company has an organization in AWS Organizations. The organization includes workload accounts that contain enterprise applications. The company centrally manages users from an operations account. No users can be created in the workload accounts. The company recently added an operations team and must provide the operations team members with administrator access to each workload account. Which combination of actions will provide this access? (Choose three.)

- A. Create a SysAdmin role in the operations account
- B. Attach the AdministratorAccess policy to the role
- C. Modify the trust relationship to allow the sts:AssumeRole action from the workload accounts.
- D. Create a SysAdmin role in each workload account
- E. Attach the AdministratorAccess policy to the role
- F. Modify the trust relationship to allow the sts:AssumeRole action from the operations account.
- G. Create an Amazon Cognito identity pool in the operations account
- H. Attach the SysAdmin role as an authenticated role.
- I. In the operations account, create an IAM user for each operations team member.
- J. In the operations account, create an IAM user group that is named SysAdmin
- K. Add an IAM policy that allows the sts:AssumeRole action for the SysAdmin role in each workload account
- L. Add all operations team members to the group.
- M. Create an Amazon Cognito user pool in the operations account
- N. Create an Amazon Cognito user for each operations team member.

Answer: BDE

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 17

A DevOps engineer is deploying a new version of a company's application in an AWS CodeDeploy deployment group associated with its Amazon EC2 instances. After some time, the deployment fails. The engineer realizes that all the events associated with the specific deployment ID are in a Skipped status and code was not deployed in the instances associated with the deployment group. What are valid reasons for this failure? (Select TWO.).

- A. The networking configuration does not allow the EC2 instances to reach the internet via a NAT gateway or internet gateway and the CodeDeploy endpoint cannot be reached.
- B. The IAM user who triggered the application deployment does not have permission to interact with the CodeDeploy endpoint.
- C. The target EC2 instances were not properly registered with the CodeDeploy endpoint.
- D. An instance profile with proper permissions was not attached to the target EC2 instances.
- E. The appspec
- F. yml file was not included in the application revision.

Answer: AD

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html#troubleshooting-skipped-lifecycle-events>

NEW QUESTION 21

A company runs a workload on Amazon EC2 instances. The company needs a control that requires the use of Instance Metadata Service Version 2 (IMDSv2) on all EC2 instances in the AWS account. If an EC2 instance does not prevent the use of Instance Metadata Service Version 1 (IMDSv1), the EC2 instance must be terminated. Which solution will meet these requirements?

- A. Set up AWS Config in the account
- B. Use a managed rule to check EC2 instance
- C. Configure the rule to remediate the findings by using AWS Systems Manager Automation to terminate the instance.
- D. Create a permissions boundary that prevents the ec2:RunInstance action if the ec2:MetadataHttpTokens condition key is not set to a value of require
- E. Attach the permissions boundary to the IAM role that was used to launch the instance.
- F. Set up Amazon Inspector in the account
- G. Configure Amazon Inspector to activate deep inspection for EC2 instance
- H. Create an Amazon EventBridge rule for an Inspector2 finding
- I. Set an AWS Lambda function as the target to terminate the instance.
- J. Create an Amazon EventBridge rule for the EC2 instance launch successful event
- K. Send the event to an AWS Lambda function to inspect the EC2 metadata and to terminate the instance.

Answer: B

Explanation:

To implement a control that requires the use of IMDSv2 on all EC2 instances in the account, the DevOps engineer can use a permissions boundary. A permissions boundary is a policy that defines the maximum permissions that an IAM entity can have. The DevOps engineer can create a permissions boundary that prevents the ec2:RunInstance action if the ec2:MetadataHttpTokens condition key is not set to a value of required. This condition key enforces the use of IMDSv2 on EC2 instances. The DevOps engineer can attach the permissions boundary to the IAM role that was used to launch the instance. This way, any attempt to launch an EC2 instance without using IMDSv2 will be denied by the permissions boundary.

NEW QUESTION 23

A company has enabled all features for its organization in AWS Organizations. The organization contains 10 AWS accounts. The company has turned on AWS CloudTrail in all the accounts. The company expects the number of AWS accounts in the organization to increase to 500 during the next year. The company plans to use multiple OUs for these accounts. The company has enabled AWS Config in each existing AWS account in the organization. A DevOps engineer must implement a solution that enables AWS Config automatically for all future AWS accounts that are created in the organization. Which solution will meet this requirement?

- A. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call.
- B. Configure the rule to invoke an AWS Lambda function that enables trusted access to AWS Config for the organization.
- C. In the organization's management account, create an AWS CloudFormation stack set to enable AWS Config.
- D. Configure the stack set to deploy automatically when an account is created through Organizations.
- E. In the organization's management account, create an SCP that allows the appropriate AWS Config API calls to enable AWS Config.
- F. Apply the SCP to the root-level OU.
- G. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call.
- H. Configure the rule to invoke an AWS Systems Manager Automation runbook to enable AWS Config for the account.

Answer: B

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2020/02/aws-cloudformation-stacksets-introduces-automatic-deployments-across-accounts-and-regions-through-aws-organizations/>

NEW QUESTION 25

A company's developers use Amazon EC2 instances as remote workstations. The company is concerned that users can create or modify EC2 security groups to allow unrestricted inbound access.

A DevOps engineer needs to develop a solution to detect when users create unrestricted security group rules. The solution must detect changes to security group rules in near real time, remove unrestricted rules, and send email notifications to the security team. The DevOps engineer has created an AWS Lambda function that checks for security group ID from input, removes rules that grant unrestricted access, and sends notifications through Amazon Simple Notification Service (Amazon SNS).

What should the DevOps engineer do next to meet the requirements?

- A. Configure the Lambda function to be invoked by the SNS topic.
- B. Create an AWS CloudTrail subscription for the SNS topic.
- C. Configure a subscription filter for security group modification events.
- D. Create an Amazon EventBridge scheduled rule to invoke the Lambda function.
- E. Define a schedule pattern that runs the Lambda function every hour.
- F. Create an Amazon EventBridge event rule that has the default event bus as the source.
- G. Define the rule's event pattern to match EC2 security group creation and modification events.
- H. Configure the rule to invoke the Lambda function.
- I. Create an Amazon EventBridge custom event bus that subscribes to events from all AWS services.
- J. Configure the Lambda function to be invoked by the custom event bus.

Answer: C

Explanation:

To meet the requirements, the DevOps engineer should create an Amazon EventBridge event rule that has the default event bus as the source. The rule's event pattern should match EC2 security group creation and modification events, and it should be configured to invoke the Lambda function. This solution will allow for near real-time detection of security group rule changes and will trigger the Lambda function to remove any unrestricted rules and send email notifications to the security team. <https://repost.aws/knowledge-center/monitor-security-group-changes-ec2>

NEW QUESTION 26

A company uses an Amazon API Gateway regional REST API to host its application API. The REST API has a custom domain. The REST API's default endpoint is deactivated.

The company's internal teams consume the API. The company wants to use mutual TLS between the API and the internal teams as an additional layer of authentication.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Use AWS Certificate Manager (ACM) to create a private certificate authority (CA). Provision a client certificate that is signed by the private CA.
- B. Provision a client certificate that is signed by a public certificate authority (CA). Import the certificate into AWS Certificate Manager (ACM).
- C. Upload the provisioned client certificate to an Amazon S3 bucket.
- D. Configure the API Gateway mutual TLS to use the client certificate that is stored in the S3 bucket as the trust store.
- E. Upload the provisioned client certificate private key to an Amazon S3 bucket.
- F. Configure the API Gateway mutual TLS to use the private key that is stored in the S3 bucket as the trust store.
- G. Upload the root private certificate authority (CA) certificate to an Amazon S3 bucket.
- H. Configure the API Gateway mutual TLS to use the private CA certificate that is stored in the S3 bucket as the trust store.

Answer: AE

Explanation:

Mutual TLS (mTLS) authentication requires two-way authentication between the client and the server. For Amazon API Gateway, you can enable mTLS for a custom domain name, which requires clients to present X.509 certificates to verify their identity to access your API. To set up mTLS, you would typically use AWS Certificate Manager (ACM) to create a private certificate authority (CA) and provision a client certificate signed by this private CA. The root CA certificate is then uploaded to an Amazon S3 bucket and configured in API Gateway as the trust store¹².

References:

- ? Introducing mutual TLS authentication for Amazon API Gateway¹.
- ? Configuring mutual TLS authentication for a REST API².
- ? AWS Private Certificate Authority details³.
- ? AWS Certificate Manager Private Certificate Authority updates⁴.

NEW QUESTION 30

A video-sharing company stores its videos in Amazon S3. The company has observed a sudden increase in video access requests, but the company does not know which videos are most popular. The company needs to identify the general access pattern for the video files. This pattern includes the number of users who access a certain file on a given day, as well as the number of times a file is accessed. A DevOps engineer manages a large commercial website that runs on Amazon EC2. The website uses Amazon Kinesis Data Streams to collect and process web logs. The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2. Sudden increases of data cause the Kinesis consumer application to fall behind and the Kinesis data streams drop records before the records can be processed. The DevOps engineer must implement a solution to improve stream handling.

Which solution meets these requirements with the MOST operational efficiency for pull requests for certain files?

How can the company meet these requirements with the LEAST amount of effort?

- A. Activate S3 server access login
- B. Import the access logs into an Amazon Aurora databas
- C. Use an Aurora SQL query to analyze the access patterns.
- D. Activate S3 server access login
- E. Use Amazon Athena to create an external table with the log file
- F. Use Athena to create a SQL query to analyze the access patterns.
- G. Invoke an AWS Lambda function for every S3 object access even
- H. Configure the Lambda function to write the file access information, such as use
- I. S3 bucket, and file key, to an Amazon Aurora databas
- J. Use an Aurora SQL query to analyze the access patterns.
- K. Record an Amazon CloudWatch Logs log message for every S3 object access even
- L. Configure a CloudWatch Logs log stream to write the file access information, such as user, S3 bucket, and file key, to an Amazon Kinesis Data Analytics for SQL applicatio
- M. Perform a sliding window analysis.

Answer: B

Explanation:

Activating S3 server access logging and using Amazon Athena to create an external table with the log files is the easiest and most cost-effective way to analyze access patterns. This option requires minimal setup and allows for quick analysis of the access patterns with SQL queries. Additionally, Amazon Athena scales automatically to match the query load, so there is no need for additional infrastructure provisioning or management.

NEW QUESTION 34

A DevOps engineer needs to back up sensitive Amazon S3 objects that are stored within an S3 bucket with a private bucket policy using S3 cross-Region replication functionality. The objects need to be copied to a target bucket in a different AWS Region and account. Which combination of actions should be performed to enable this replication? (Choose three.)

- A. Create a replication IAM role in the source account
- B. Create a replication I AM role in the target account.
- C. Add statements to the source bucket policy allowing the replication IAM role to replicate objects.
- D. Add statements to the target bucket policy allowing the replication IAM role to replicate objects.
- E. Create a replication rule in the source bucket to enable the replication.
- F. Create a replication rule in the target bucket to enable the replication.

Answer: ADE

Explanation:

S3 cross-Region replication (CRR) automatically replicates data between buckets across different AWS Regions. To enable CRR, you need to add a replication configuration to your source bucket that specifies the destination bucket, the IAM role, and the encryption type (optional). You also need to grant permissions to the IAM role to perform replication actions on both the source and destination buckets. Additionally, you can choose the destination storage class and enable additional replication options such as S3 Replication Time Control (S3 RTC) or S3 Batch Replication. <https://medium.com/cloud-techies/s3-same-region-replication-srr-and-cross-region-replication-crr-34d446806bab> <https://aws.amazon.com/getting-started/hands-on/replicate-data-using-amazon-s3-replication/> <https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html>

NEW QUESTION 36

A company has multiple AWS accounts. The company uses AWS IAM Identity Center (AWS Single Sign-On) that is integrated with AWS Toolkit for Microsoft Azure DevOps. The attributes for access control feature is enabled in IAM Identity Center.

The attribute mapping list contains two entries. The department key is mapped to

`${path:enterprise.department}`. The costCenter key is mapped to

`${path:enterprise.costCenter}`.

All existing Amazon EC2 instances have a department tag that corresponds to three company departments (d1, d2, d3). A DevOps engineer must create policies based on the matching attributes. The policies must minimize administrative effort and must grant each Azure AD user access to only the EC2 instances that are tagged with the user's respective department name.

Which condition key should the DevOps engineer include in the custom permissions policies to meet these requirements?

A.

```
"Condition": {
  "ForAllValues:StringEquals": {
    "aws:TagKeys": ["department"]
  }
}
```

B.

```
"Condition": {
  "StringEquals": {
    "aws:PrincipalTag/department": "${aws:ResourceTag/department}"
  }
}
```

C.

```
"Condition": {
  "StringEquals": {
    "ec2:ResourceTag/department": "${aws:PrincipalTag/department}"
  }
}
```

D.

```
"Condition": {
  "ForAllValues:StringEquals": {
    "ec2:ResourceTag/department": ["d1", "d2", "d3"]
  }
}
```

A.

Answer: C

Explanation:

<https://docs.aws.amazon.com/singlesignon/latest/userguide/configure-abac.html>

NEW QUESTION 37

A company recently migrated its legacy application from on-premises to AWS. The application is hosted on Amazon EC2 instances behind an Application Load Balancer which is behind Amazon API Gateway. The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application. The company also wants to ensure it can quickly roll back updates if there is an issue. Which solution will meet these requirements with MINIMAL changes to the application?

- A. Introduce changes as a separate environment parallel to the existing one Configure API Gateway to use a canary release deployment to send a small subset of user traffic to the new environment.
- B. Introduce changes as a separate environment parallel to the existing one Update the application's DNS alias records to point to the new environment.
- C. Introduce changes as a separate target group behind the existing Application Load Balancer Configure API Gateway to route user traffic to the new target group in steps.
- D. Introduce changes as a separate target group behind the existing Application Load Balancer Configure API Gateway to route all traffic to the Application Load Balancer which then sends the traffic to the new target group.

Answer: A

Explanation:

API Gateway supports canary deployment on a deployment stage before you direct all traffic to that stage. A parallel environment means we will create a new ALB and a target group that will target a new set of EC2 instances on which the newer version of the app will be deployed. So the canary setting associated to the new version of the API will connect with the new ALB instance which in turn will direct the traffic to the new EC2 instances on which the newer version of the application is deployed.

NEW QUESTION 42

A company's application development team uses Linux-based Amazon EC2 instances as bastion hosts. Inbound SSH access to the bastion hosts is restricted to specific IP addresses, as defined in the associated security groups. The company's security team wants to receive a notification if the security group rules are modified to allow SSH access from any IP address. What should a DevOps engineer do to meet this requirement?

- A. Create an Amazon EventBridge rule with a source of aws.cloudtrail and the event name AuthorizeSecurityGroupIngres
- B. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
- C. Enable Amazon GuardDuty and check the findings for security groups in AWS Security Hub
- D. Configure an Amazon EventBridge rule with a custom pattern that matches GuardDuty events with an output of NON_COMPLIAN
- E. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
- F. Create an AWS Config rule by using the restricted-ssh managed rule to check whether security groups disallow unrestricted incoming SSH traffic
- G. Configure automatic remediation to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
- H. Enable Amazon Inspector
- I. Include the Common Vulnerabilities and Exposures-1.1 rules package to check the security groups that are associated with the bastion host
- J. Configure Amazon Inspector to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/monitor-security-group-changes-ec2/>

NEW QUESTION 43

A company has many applications. Different teams in the company developed the applications by using multiple languages and frameworks. The applications run on premises and on different servers with different operating systems. Each team has its own release protocol and process. The company wants to reduce the complexity of the release and maintenance of these applications.

The company is migrating its technology stacks, including these applications, to AWS. The

company wants centralized control of source code, a consistent and automatic delivery pipeline, and as few maintenance tasks as possible on the underlying infrastructure.

What should a DevOps engineer do to meet these requirements?

- A. Create one AWS CodeCommit repository for all application
- B. Put each application's code in a different branch
- C. Merge the branches, and use AWS CodeBuild to build the application
- D. Use AWS CodeDeploy to deploy the applications to one centralized application server.
- E. Create one AWS CodeCommit repository for each of the application
- F. Use AWS CodeBuild to build the applications one at a time
- G. Use AWS CodeDeploy to deploy the applications to one centralized application server.
- H. Create one AWS CodeCommit repository for each of the application
- I. Use AWS CodeBuild to build the applications one at a time and to create one AMI for each server
- J. Use AWS CloudFormation StackSets to automatically provision and decommission Amazon EC2 fleets by using these AMIs.
- K. Create one AWS CodeCommit repository for each of the application
- L. Use AWS CodeBuild to build one Docker image for each application in Amazon Elastic Container Registry (Amazon ECR). Use AWS CodeDeploy to deploy the applications to Amazon Elastic Container Service (Amazon ECS) on infrastructure that AWS Fargate manages.

Answer: D

Explanation:

because of "as few maintenance tasks as possible on the underlying infrastructure". Fargate does that better than "one centralized application server"

NEW QUESTION 48

A company provides an application to customers. The application has an Amazon API Gateway REST API that invokes an AWS Lambda function. On initialization, the Lambda function loads a large amount of data from an Amazon DynamoDB table. The data load process results in long cold-start times of 8-10 seconds. The DynamoDB table has DynamoDB Accelerator (DAX) configured.

Customers report that the application intermittently takes a long time to respond to requests. The application receives thousands of requests throughout the day. In the middle of the day, the application experiences 10 times more requests than at any other time of the day. Near the end of the day, the application's request volume decreases to 10% of its normal total.

A DevOps engineer needs to reduce the latency of the Lambda function at all times of the day.

Which solution will meet these requirements?

- A. Configure provisioned concurrency on the Lambda function with a concurrency value of 1. Delete the DAX cluster for the DynamoDB table.
- B. Configure reserved concurrency on the Lambda function with a concurrency value of 0.
- C. Configure provisioned concurrency on the Lambda function
- D. Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.
- E. Configure reserved concurrency on the Lambda function
- F. Configure AWS Application Auto Scaling on the API Gateway API with a reserved concurrency maximum value of 100.

Answer: C

Explanation:

The following are the steps that the DevOps engineer should take to reduce the latency of the Lambda function at all times of the day:

? Configure provisioned concurrency on the Lambda function.

? Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.

The provisioned concurrency setting ensures that there is always a minimum number of Lambda function instances available to handle requests. The Application Auto Scaling setting will automatically scale the number of Lambda function instances up or down based on the demand for the application.

This solution will ensure that the Lambda function is able to handle the increased load during the middle of the day, while also keeping the cold-start latency low.

The following are the reasons why the other options are not correct:

? Option A is incorrect because it will not reduce the cold-start latency of the Lambda function.

? Option B is incorrect because it will not scale the number of Lambda function instances up or down based on demand.

? Option D is incorrect because it will only configure reserved concurrency on the API Gateway API, which will not affect the Lambda function.

NEW QUESTION 53

A company's application teams use AWS CodeCommit repositories for their applications.

The application teams have repositories in multiple AWS accounts. All accounts are in an organization in AWS Organizations.

Each application team uses AWS IAM Identity Center (AWS Single Sign-On) configured with an external IdP to assume a developer IAM role. The developer role allows the application teams to use Git to work with the code in the repositories.

A security audit reveals that the application teams can modify the main branch in any repository. A DevOps engineer must implement a solution that allows the application teams to modify the main branch of only the repositories that they manage.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Update the SAML assertion to pass the user's team name
- B. Update the IAM role's trust policy to add an access-team session tag that has the team name.
- C. Create an approval rule template for each team in the Organizations management account
- D. Associate the template with all the repositories
- E. Add the developer role ARN as an approver.
- F. Create an approval rule template for each account
- G. Associate the template with all repositories
- H. Add the "aws:ResourceTag/access-team":"\$;{aws:PrincipalTag/access-team}" condition to the approval rule template.
- I. For each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.
- J. Attach an SCP to the account
- K. Include the following statement:


```
{
  "Effect": "Deny",
  "Action": [
    "codecommit:GitPush",
    "codecommit:PutFile",
    "codecommit:Merge*"
  ],
  "Resource": "*",
  "Condition": {
    "StringEqualsIfExists": {
      "codecommit:References": ["refs/heads/main"]
    },
    "StringNotEquals": {
      "aws:ResourceTag/access-team": "$ ;{aws:PrincipalTag/access-team}"
    },
    "Null": {
      "codecommit:References": "false"
    }
  }
}
```

L. Create an IAM permissions boundary in each account

M. Include the following statement: {

```
  "Effect": "Allow",
  "Action": [
    "codecommit:GitPush",
    "codecommit:PutFile",
    "codecommit:Merge*"
  ],
  "Resource": "*",
  "Condition": {
    "StringEqualsIfExists": {
      "codecommit:References": ["refs/heads/main"]
    },
    "StringNotEquals": {
      "aws:ResourceTag/access-team": "$ ;{aws:PrincipalTag/access-team}"
    },
    "Null": {
      "codecommit:References": "false"
    }
  }
}
```

Answer: ADF

Explanation:

Short Explanation: To meet the requirements, the DevOps engineer should update the SAML assertion to pass the user's team name, update the IAM role's trust policy to add an access-team session tag that has the team name, create an IAM permissions boundary in each account, and for each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.

References:

? Updating the SAML assertion to pass the user's team name allows the DevOps engineer to use IAM tags to identify which team a user belongs to. This can help enforce fine-grained access control based on the user's team membership1.

? Updating the IAM role's trust policy to add an access-team session tag that has the team name allows the DevOps engineer to use IAM condition keys to restrict access based on the session tag value2. For example, the DevOps engineer can use the aws:PrincipalTag condition key to match the access-team tag of the user with the access-team tag of the repository3.

? Creating an IAM permissions boundary in each account allows the DevOps engineer to set the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries4. For example, the DevOps engineer can use a permissions boundary policy to limit the actions that a user can perform on CodeCommit repositories based on their access-team tag5.

? For each CodeCommit repository, adding an access-team tag that has the value set to the name of the associated team allows the DevOps engineer to use resource tags to identify which team manages a repository. This can help enforce fine-grained access control based on the resource tag value6.

? The other options are incorrect because:

NEW QUESTION 56

A company has 20 service teams. Each service team is responsible for its own microservice. Each service team uses a separate AWS account for its microservice and a VPC with the 192.168.0.0/22 CIDR block. The company manages the AWS accounts with AWS Organizations.

Each service team hosts its microservice on multiple Amazon EC2 instances behind an Application Load Balancer. The microservices communicate with each other across the public internet. The company's security team has issued a new guideline that all communication between microservices must use HTTPS over private network connections and cannot traverse the public internet.

A DevOps engineer must implement a solution that fulfills these obligations and minimizes the number of changes for each service team.

Which solution will meet these requirements?

A. Create a new AWS account in AWS Organizations. Create a VPC in this account and use AWS Resource Access Manager to share the private subnets of this

VPC with the organization Instruct the service teams to launch a ne

- B. Network Load Balancer (NLB) and EC2 instances that use the shared private subnets Use the NLB DNS names for communication between microservices.
- C. Create a Network Load Balancer (NLB) in each of the microservice VPCs Use AWS PrivateLink to create VPC endpoints in each AWS account for the NLBs Create subscriptions to each VPC endpoint in each of the other AWS accounts Use the VPC endpoint DNS names for communication between microservices.
- D. Create a Network Load Balancer (NLB) in each of the microservice VPCs Create VPC peering connections between each of the microservice VPCs Update the route tables for each VPC to use the peering links Use the NLB DNS names for communication between microservices.
- E. Create a new AWS account in AWS Organizations Create a transit gateway in this account and use AWS Resource Access Manager to share the transit gateway with the organizatio
- F. In each of the microservice VPC
- G. create a transit gateway attachment to the shared transit gateway Update the route tables of each VPC to use the transit gateway Create a Network Load Balancer (NLB) in each of the microservice VPCs Use the NLB DNS names for communication between microservices.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/> Private link is the best option because Transit Gateway doesn't support overlapping CIDR ranges.

NEW QUESTION 61

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function that uses reserved concurrency. The Lambda function processes order messages from an Amazon Simple Queue Service (Amazon SQS) queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has auto scaling enabled for read and write capacity.

Which actions should a DevOps engineer take to resolve this delay? (Choose two.)

- A. Check the ApproximateAgeOfOldestMessage metric for the SQS queu
- B. Increase the Lambda function concurrency limit.
- C. Check the ApproximateAgeOfOldestMessage metnc for the SQS queue Configure a redrive policy on the SQS queue.
- D. Check the NumberOfMessagesSent metric for the SQS queu
- E. Increase the SQS queue visibility timeout.
- F. Check the WriteThrottleEvents metric for the DynamoDB tabl
- G. Increase the maximum write capacity units (WCUs) for the table's scaling policy.
- H. Check the Throttles metric for the Lambda functio
- I. Increase the Lambda function timeout.

Answer: AD

Explanation:

A: If the ApproximateAgeOfOldestMessages indicate that orders are remaining in the SQS queue for longer than expected, the reserved concurrency limit may be set too small to keep up with the number of orders entering the queue and is being throttled. D: The DynamoDB table is using Auto Scaling. With Auto Scaling, you create a scaling policy that specifies whether you want to scale read capacity or write capacity (or both), and the minimum and maximum provisioned capacity unit settings for the table. The ThottledWriteRequests metric will indicate if there is a throttling issue on the DynamoDB table, which can be resolved by increasing the maximum write capacity units for the table's Auto Scaling policy. <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

NEW QUESTION 62

An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages while it is running.

All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.

How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

- A. Add a DelelionPolicy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.
- B. Add a custom resource with an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM rol
- C. Write the Lambda function to delete all objects from the bucket when RequestType is Delete.
- D. Identify the resource that was not delete
- E. Manually empty the S3 bucket and then delete it.
- F. Replace the EC2 and S3 bucket resources with a single AWS OpsWorks Stacks resourc
- G. Define a custom recipe for the stack to create and delete the EC2 instance and the S3 bucket.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-s3-custom-resources/>

NEW QUESTION 66

A company runs its container workloads in AWS App Runner. A DevOps engineer manages the company's container repository in Amazon Elastic Container Registry (Amazon ECR).

The DevOps engineer must implement a solution that continuously monitors the container repository. The solution must create a new container image when the solution detects an operating system vulnerability or language package vulnerability.

Which solution will meet these requirements?

- A. Use EC2 Image Builder to create a container image pipelin
- B. Use Amazon ECR as the target repositor
- C. Turn on enhanced scanning on the ECR repositor
- D. Create an Amazon EventBridge rule to capture an Inspector2 finding even
- E. Use the event to invoke the image pipelin
- F. Re-upload the container to the repository.
- G. Use EC2 Image Builder to create a container image pipelin
- H. Use Amazon ECR as the target repositor

- I. Enable Amazon GuardDuty Malware Protection on the container workloa
- J. Create an Amazon EventBridge rule to capture a GuardDuty finding even
- K. Use the event to invoke the image pipeline.
- L. Create an AWS CodeBuild project to create a container imag
- M. Use Amazon ECR as the target repositor
- N. Turn on basic scanning on the repositor
- O. Create an Amazon EventBridge rule to capture an ECR image action even
- P. Use the event to invoke the CodeBuild projec
- Q. Re-upload the container to the repository.
- R. Create an AWS CodeBuild project to create a container imag
- S. Use Amazon ECR as the target repositor
- T. Configure AWS Systems Manager Compliance to scan all managed node
- . Create an Amazon EventBridge rule to capture a configuration compliance state change even
- . Use the event to invoke the CodeBuild project.

Answer: A

Explanation:

The solution that meets the requirements is to use EC2 Image Builder to create a container image pipeline, use Amazon ECR as the target repository, turn on enhanced scanning on the ECR repository, create an Amazon EventBridge rule to capture an Inspector2 finding event, and use the event to invoke the image pipeline. Re-upload the container to the repository.

This solution will continuously monitor the container repository for vulnerabilities using enhanced scanning, which is a feature of Amazon ECR that provides detailed information and guidance on how to fix security issues found in your container images. Enhanced scanning uses Inspector2, a security assessment service that integrates with Amazon ECR and generates findings for any vulnerabilities detected in your images. You can use Amazon EventBridge to create a rule that triggers an action when an Inspector2 finding event occurs. The action can be to invoke an EC2 Image Builder pipeline, which is a service that automates the creation of container images. The pipeline can use the latest patches and updates to build a new container image and upload it to the same ECR repository, replacing the vulnerable image.

The other options are not correct because they do not meet all the requirements or use services that are not relevant for the scenario.

Option B is not correct because it uses Amazon GuardDuty Malware Protection, which is a feature of GuardDuty that detects malicious activity and unauthorized behavior on your AWS accounts and resources. GuardDuty does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder.

Option C is not correct because it uses basic scanning on the ECR repository, which only provides a summary of the vulnerabilities found in your container images. Basic scanning does not use Inspector2 or generate findings that can be captured by Amazon EventBridge. Moreover, basic scanning does not provide guidance on how to fix the vulnerabilities.

Option D is not correct because it uses AWS Systems Manager Compliance, which is a feature of Systems Manager that helps you monitor and manage the compliance status of your AWS resources based on AWS Config rules and AWS Security Hub standards. Systems Manager Compliance does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder.

NEW QUESTION 70

A company hosts applications in its AWS account Each application logs to an individual Amazon CloudWatch log group. The company's CloudWatch costs for ingestion are increasing

A DevOps engineer needs to Identify which applications are the source of the increased logging costs.

Which solution Will meet these requirements?

- A. Use CloudWatch metrics to create a custom expression that Identifies the CloudWatch log groups that have the most data being written to them.
- B. Use CloudWatch Logs Insights to create a set of queries for the application log groups to Identify the number of logs written for a period of time
- C. Use AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage
- D. Use AWS CloudTrail to filter for CreateLogStream events for each application

Answer: C

Explanation:

The correct answer is C.

A comprehensive and detailed explanation is:

? Option A is incorrect because using CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them is not a valid solution. CloudWatch metrics do not provide information about the size or volume of data being ingested by CloudWatch logs.

CloudWatch metrics only provide information about the number of events, bytes, and errors that occur within a log group or stream. Moreover, creating a custom expression with CloudWatch metrics would require using the search_web tool, which is not necessary for this use case.

? Option B is incorrect because using CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time is not a valid solution. CloudWatch Logs Insights can help analyze and filter log events based on patterns and expressions, but it does not provide information about the cost or billing of CloudWatch logs. CloudWatch Logs Insights also charges based on the amount of data scanned by each query, which could increase the logging costs further.

? Option C is correct because using AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage is a valid solution. AWS Cost Explorer is a tool that helps visualize, understand, and manage AWS costs and usage over time. AWS Cost Explorer can generate custom reports that show the breakdown of costs by service, region, account, tag, or any other dimension. AWS Cost Explorer can also filter and group costs by usage type, which can help identify the specific CloudWatch log groups that are the source of the increased logging costs.

? Option D is incorrect because using AWS CloudTrail to filter for CreateLogStream events for each application is not a valid solution. AWS CloudTrail is a service that records API calls and account activity for AWS services, including CloudWatch logs. However, AWS CloudTrail does not provide information about the cost or billing of CloudWatch logs. Filtering for CreateLogStream events would only show when a new log stream was created within a log group, but not how much data was ingested or stored by that log stream.

References:

? CloudWatch Metrics

? CloudWatch Logs Insights

? AWS Cost Explorer

? AWS CloudTrail

NEW QUESTION 73

A company has its AWS accounts in an organization in AWS Organizations. AWS Config is manually configured in each AWS account. The company needs to implement a solution to centrally configure AWS Config for all accounts in the organization The solution also must record resource changes to a central account.

Which combination of actions should a DevOps engineer perform to meet these requirements? (Choose two.)

- A. Configure a delegated administrator account for AWS Config
- B. Enable trusted access for AWS Config in the organization.
- C. Configure a delegated administrator account for AWS Config
- D. Create a service-linked role for AWS Config in the organization's management account.
- E. Create an AWS CloudFormation template to create an AWS Config aggregator
- F. Configure a CloudFormation stack set to deploy the template to all accounts in the organization.
- G. Create an AWS Config organization aggregator in the organization's management account
- H. Configure data collection from all AWS accounts in the organization and from all AWS Regions.
- I. Create an AWS Config organization aggregator in the delegated administrator account
- J. Configure data collection from all AWS accounts in the organization and from all AWS Regions.

Answer: AE

Explanation:

<https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/> <https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html>

NEW QUESTION 74

A company needs to ensure that flow logs remain configured for all existing and new VPCs in its AWS account. The company uses an AWS CloudFormation stack to manage its VPCs. The company needs a solution that will work for any VPCs that any IAM user creates. Which solution will meet these requirements?

- A. Add the resource to the CloudFormation stack that creates the VPCs.
- B. Create an organization in AWS Organization
- C. Add the company's AWS account to the organization
- D. Create an SCP to prevent users from modifying VPC flow logs.
- E. Turn on AWS Config
- F. Create an AWS Config rule to check whether VPC flow logs are turned on
- G. Configure automatic remediation to turn on VPC flow logs.
- H. Create an IAM policy to deny the use of API calls for VPC flow logs
- I. Attach the IAM policy to all IAM users.

Answer: C

Explanation:

To meet the requirements of ensuring that flow logs remain configured for all existing and new VPCs in the AWS account, the company should use AWS Config and automatic remediation. AWS Config is a service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records the configuration changes of the AWS resources and evaluates them against desired configurations. Customers can use AWS Config rules to define the desired configuration state of their AWS resources and trigger actions when a resource configuration violates a rule.

One of the AWS Config rules that customers can use is `vpc-flow-logs-enabled`, which checks whether VPC flow logs are enabled for all VPCs in an AWS account. Customers can also configure automatic remediation for this rule, which means that AWS Config will automatically enable VPC flow logs for any VPCs that do not have them enabled. Customers can specify the destination (CloudWatch Logs or S3) and the traffic type (all, accept, or reject) for the flow logs as remediation parameters. By using AWS Config and automatic remediation, the company can ensure that flow logs remain configured for all existing and new VPCs in its AWS account, regardless of who creates them or how they are created.

The other options are not correct because they do not meet the requirements or follow best practices. Adding the resource to the CloudFormation stack that creates the VPCs is not a sufficient solution because it will only work for VPCs that are created by using the CloudFormation stack. It will not work for VPCs that are created by using other methods, such as the console or the API. Creating an organization in AWS Organizations and creating an SCP to prevent users from modifying VPC flow logs is not a good solution because it will not ensure that flow logs are enabled for all VPCs in the first place. It will only prevent users from disabling or changing flow logs after they are enabled. Creating an IAM policy to deny the use of API calls for VPC flow logs and attaching it to all IAM users is not a valid solution because it will prevent users from enabling or disabling flow logs at all.

It will also not work for VPCs that are created by using other methods, such as the console or CloudFormation.

References:

- ? 1: `AWS::EC2::FlowLog` - AWS CloudFormation
- ? 2: Amazon VPC Flow Logs extends CloudFormation Support to custom format subscriptions, 1-minute aggregation intervals and tagging
- ? 3: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud
- ? : About AWS Config - AWS Config
- ? : `vpc-flow-logs-enabled` - AWS Config
- ? : Remediate Noncompliant Resources with AWS Config Rules - AWS Config

NEW QUESTION 79

A company is deploying a new application that uses Amazon EC2 instances. The company needs a solution to query application logs and AWS account API activity. Which solution will meet these requirements?

- A. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs. Configure AWS CloudTrail to deliver the API logs to Amazon S3. Use CloudWatch to query both sets of logs.
- B. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs. Configure AWS CloudTrail to deliver the API logs to CloudWatch Logs. Use CloudWatch Logs Insights to query both sets of logs.
- C. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon Kinesis. Configure AWS CloudTrail to deliver the API logs to Kinesis. Use Kinesis to load the data into Amazon Redshift. Use Amazon Redshift to query both sets of logs.
- D. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon S3. Use AWS CloudTrail to deliver the API logs to Amazon S3. Use Amazon Athena to query both sets of logs in Amazon S3.

Answer: D

Explanation:

This solution will meet the requirements because it will use Amazon S3 as a common data lake for both the application logs and the API logs. Amazon S3 is a service that provides scalable, durable, and secure object storage for any type of data. You can use the Amazon CloudWatch agent to send logs from your EC2 instances to S3 buckets, and use AWS CloudTrail to deliver the API logs to S3 buckets as well. You can also use Amazon Athena to query both sets of logs in S3 using standard SQL, without loading or transforming them. Athena is a serverless interactive query service that allows you to analyze data in S3 using a variety of data formats, such as JSON, CSV, Parquet, and ORC.

NEW QUESTION 80

A company is performing vulnerability scanning for all Amazon EC2 instances across many accounts. The accounts are in an organization in AWS Organizations. Each account's VPCs are attached to a shared transit gateway. The VPCs send traffic to the internet through a central egress VPC. The company has enabled Amazon Inspector in a delegated administrator account and has enabled scanning for all member accounts. A DevOps engineer discovers that some EC2 instances are listed in the "not scanning" tab in Amazon Inspector. Which combination of actions should the DevOps engineer take to resolve this issue? (Choose three.)

- A. Verify that AWS Systems Manager Agent is installed and is running on the EC2 instances that Amazon Inspector is not scanning.
- B. Associate the target EC2 instances with security groups that allow outbound communication on port 443 to the AWS Systems Manager service endpoint.
- C. Grant inspector: StartAssessmentRun permissions to the IAM role that the DevOps engineer is using.
- D. Configure EC2 Instance Connect for the EC2 instances that Amazon Inspector is not scanning.
- E. Associate the target EC2 instances with instance profiles that grant permissions to communicate with AWS Systems Manager.
- F. Create a managed-instance activation
- G. Use the Activation Code and the Activation ID to register the EC2 instances.

Answer: ABE

Explanation:

<https://docs.aws.amazon.com/inspector/latest/user/scanning-ec2.html>

NEW QUESTION 83

A company uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups. The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account. When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account's backup vault. Which combination of steps must the company take so that backups can be copied to the new account's backup vault? (Select TWO.)

- A. Edit the backup vault access policy in the new account to allow access to the primary account.
- B. Edit the backup vault access policy in the primary account to allow access to the new account.
- C. Edit the backup vault access policy in the primary account to allow access to the KMS key in the new account.
- D. Edit the key policy of the KMS key in the primary account to share the key with the new account.
- E. Edit the key policy of the KMS key in the new account to share the key with the primary account.

Answer: AE

Explanation:

To enable cross-account backup, the company needs to grant permissions to both the backup vault and the KMS key in the destination account. The backup vault access policy in the destination account must allow the primary account to copy backups into the vault. The key policy of the KMS key in the destination account must allow the primary account to use the key to encrypt and decrypt the backups. These steps are described in the AWS documentation¹². Therefore, the correct answer is A and E.

References:

? 1: Creating backup copies across AWS accounts - AWS Backup

? 2: Using AWS Backup with AWS Organizations - AWS Backup

NEW QUESTION 85

A highly regulated company has a policy that DevOps engineers should not log in to their Amazon EC2 instances except in emergencies. If a DevOps engineer does log in the security team must be notified within 15 minutes of the occurrence. Which solution will meet these requirements'?

- A. Install the Amazon Inspector agent on each EC2 instance Subscribe to Amazon EventBridge notifications Invoke an AWS Lambda function to check if a message is about user logins If it is send a notification to the security team using Amazon SNS.
- B. Install the Amazon CloudWatch agent on each EC2 instance Configure the agent to push all logs to Amazon CloudWatch Logs and set up a CloudWatch metric filter that searches for user login
- C. If a login is found send a notification to the security team using Amazon SNS.
- D. Set up AWS CloudTrail with Amazon CloudWatch Log
- E. Subscribe CloudWatch Logs to Amazon Kinesis Attach AWS Lambda to Kinesis to parse and determine if a log contains a user login If it does, send a notification to the security team using Amazon SNS.
- F. Set up a script on each Amazon EC2 instance to push all logs to Amazon S3 Set up an S3 event to invoke an AWS Lambda function which invokes an Amazon Athena query to ru
- G. The Athena query checks tor logins and sends the output to the security team using Amazon SNS.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/security/how-to-monitor-and-visualize-failed-ssh-access-attempts-to-amazon-ec2-linux-instances/>

NEW QUESTION 86

An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to external identity provider (IdP) and has configured SAML 2.0. The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources. Which combination of steps will meet these requirements? (Choose three.)

- A. Create IAM policies that include the required permission
- B. Include the aws:PrincipalTag condition key.
- C. Create permission set

- D. Attach an inline policy that includes the required permissions and uses the aws:PrincipalTag condition key to scope the permissions.
- E. Create a group in the Id
- F. Place users in the grou
- G. Assign the group to accounts and the permission sets in IAM Identity Center.
- H. Create a group in the Id
- I. Place users in the grou
- J. Assign the group to OUs and IAM policies.
- K. Enable attributes for access control in IAM Identity Cente
- L. Apply tags to user
- M. Map the tags as key-value pairs.
- N. Enable attributes for access control in IAM Identity Cente
- O. Map attributes from the IdP as key-value pairs.

Answer: BCF

Explanation:

Using the principalTag in the Permission Set inline policy a logged in user belonging to a specific AD group in the IDP can be permitted access to perform operations on certain resources if their group matches the group used in the PrincipleTag. Basically you are narrowing the scope of privileges assigned via Permission policies conditionally based on whether the logged in user belongs to a specific AD Group in IDP. The mapping of the AD group to the request attributes can be done using SSO attributes where we can pass other attributes like the SAML token as well.
<https://docs.aws.amazon.com/singlesignon/latest/userguide/abac.html>

NEW QUESTION 91

A company wants to use a grid system for a proprietary enterprise m-memory data store on top of AWS. This system can run in multiple server nodes in any Linux-based distribution. The system must be able to reconfigure the entire cluster every time a node is added or removed. When adding or removing nodes an /etc./cluster/nodes config file must be updated listing the IP addresses of the current node members of that cluster.

The company wants to automate the task of adding new nodes to a cluster. What can a DevOps engineer do to meet these requirements?

- A. Use AWS OpsWorks Stacks to layer the server nodes of that cluste
- B. Create a Chef recipe that populates the content of the 'etc./cluster/nodes config file and restarts the service by using the current members of the laye
- C. Assign that recipe to the Configure lifecycle event.
- D. Put the file nodes config in version contro
- E. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for thecluster node
- F. When adding a new node to the cluster update the file with all tagged instances and make a commit in version contro
- G. Deploy the new file and restart the services.
- H. Create an Amazon S3 bucket and upload a version of the /etc./cluster/nodes config file Create a crontab script that will poll for that S3 file and download it frequentl
- I. Use a process manager such as Monit or system, to restart the cluster services when it detects that the new file was modifie
- J. When adding a node to the cluster edit the file's most recent members Upload the new file to the S3 bucket.
- K. Create a user data script that lists all members of the current security group of the cluster and automatically updates the /etc/cluster/. nodes confi
- L. Tile whenever a new instance is added to the cluster.

Answer: A

Explanation:

You can run custom recipes manually, but the best approach is usually to have AWS OpsWorks Stacks run them automatically. Every layer has a set of built-in recipes assigned each of five lifecycle events—Setup, Configure, Deploy, Undeploy, and Shutdown. Each time an event occurs for an instance, AWS OpsWorks Stacks runs the associated recipes for each of the instance's layers, which handle the corresponding tasks. For example, when an instance finishes booting, AWS OpsWorks Stacks triggers a Setup event. This event runs the associated layer's Setup recipes, which typically handle tasks such as installing and configuring packages

NEW QUESTION 92

A development team wants to use AWS CloudFormation stacks to deploy an application. However, the developer IAM role does not have the required permissions to provision the resources that are specified in the AWS CloudFormation template. A DevOps engineer needs to implement a solution that allows the developers to deploy the stacks. The solution must follow the principle of least privilege.

Which solution will meet these requirements?

- A. Create an IAM policy that allows the developers to provision the required resource
- B. Attach the policy to the developer IAM role.
- C. Create an IAM policy that allows full access to AWS CloudFormatio
- D. Attach the policy to the developer IAM role.
- E. Create an AWS CloudFormation service role that has the required permission
- F. Grant the developer IAM role a cloudformation:* actio
- G. Use the new service role during stack deployments.
- H. Create an AWS CloudFormation service role that has the required permission
- I. Grant the developer IAM role the iam:PassRole permissio
- J. Use the new service role during stack deployments.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html>

NEW QUESTION 97

AnyCompany is using AWS Organizations to create and manage multiple AWS accounts AnyCompany recently acquired a smaller company, Example Corp. During the acquisition process, Example Corp's single AWS account joined AnyCompany's management account through an Organizations invitation.

AnyCompany moved the new member account under an OU that is dedicated to Example Corp.

AnyCompany's DevOps eng•neer has an IAM user that assumes a role that is named OrganizationAccountAccessRole to access member accounts. This role is configured with a full access policy When the DevOps engineer tries to use the AWS Management Console to assume the role in Example Corp's new member account, the DevOps engineer receives the following error message "Invalid information in one or more fields. Check your information or contact your

administrator."

Which solution will give the DevOps engineer access to the new member account?

- A. In the management account, grant the DevOps engineer's IAM user permission to assume the OrganizationAccountAccessRole IAM role in the new member account.
- B. In the management account, create a new SCP. In the SCP, grant the DevOps engineer's IAM user full access to all resources in the new member account.
- C. Attach the SCP to the OU that contains the new member account.
- D. In the new member account, create a new IAM role that is named OrganizationAccountAccessRole.
- E. Attach the AdministratorAccess AWS managed policy to the role.
- F. In the role's trust policy, grant the management account permission to assume the role.
- G. In the new member account edit the trust policy for the OrganizationAccountAccessRole IAM role.
- H. Grant the management account permission to assume the role.

Answer: C

Explanation:

The problem is that the DevOps engineer cannot assume the OrganizationAccountAccessRole IAM role in the new member account that joined AnyCompany's management account through an Organizations invitation. The solution is to create a new IAM role with the same name and trust policy in the new member account.

? Option A is incorrect, as it does not address the root cause of the error. The DevOps engineer's IAM user already has permission to assume the OrganizationAccountAccessRole IAM role in any member account, as this is the default role name that AWS Organizations creates when a new account joins an organization. The error occurs because the new member account does not have this role, as it was not created by AWS Organizations.

? Option B is incorrect, as it does not address the root cause of the error. An SCP is a policy that defines the maximum permissions for account members of an organization or organizational unit (OU). An SCP does not grant permissions to IAM users or roles, but rather limits the permissions that identity-based policies or resource-based policies grant to them. An SCP also does not affect how IAM roles are assumed by other principals.

? Option C is correct, as it addresses the root cause of the error. By creating a new IAM role with the same name and trust policy as the OrganizationAccountAccessRole IAM role in the new member account, the DevOps engineer can assume this role and access the account. The new role should have the AdministratorAccess AWS managed policy attached, which grants full access to all AWS resources in the account. The trust policy should allow the management account to assume the role, which can be done by specifying the management account ID as a principal in the policy statement.

? Option D is incorrect, as it assumes that the new member account already has the OrganizationAccountAccessRole IAM role, which is not true. The new member account does not have this role, as it was not created by AWS Organizations. Editing the trust policy of a non-existent role will not solve the problem.

NEW QUESTION 100

A company must encrypt all AMIs that the company shares across accounts. A DevOps engineer has access to a source account where an unencrypted custom AMI has been built. The DevOps engineer also has access to a target account where an Amazon EC2 Auto Scaling group will launch EC2 instances from the AMI. The DevOps engineer must share the AMI with the target account.

The company has created an AWS Key Management Service (AWS KMS) key in the source account.

Which additional steps should the DevOps engineer perform to meet the requirements? (Choose three.)

- A. In the source account, copy the unencrypted AMI to an encrypted AMI.
- B. Specify the KMS key in the copy action.
- C. In the source account, copy the unencrypted AMI to an encrypted AMI.
- D. Specify the default Amazon Elastic Block Store (Amazon EBS) encryption key in the copy action.
- E. In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
- F. In the source account, modify the key policy to give the target account permissions to create a grant.
- G. In the target account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role.
- H. In the source account, share the unencrypted AMI with the target account.
- I. In the source account, share the encrypted AMI with the target account.

Answer: ADF

Explanation:

The Auto Scaling group service-linked role must have a specific grant in the source account in order to decrypt the encrypted AMI. This is because the service-linked role does not have permissions to assume the default IAM role in the source account. The following steps are required to meet the requirements:

? In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.

? In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.

? In the source account, share the encrypted AMI with the target account.

? In the target account, attach the KMS grant to the Auto Scaling group service-linked role.

The first three steps are the same as the steps that I described earlier. The fourth step is required to grant the Auto Scaling group service-linked role permissions to decrypt the AMI

in the target account.

NEW QUESTION 102

A DevOps engineer is architecting a continuous development strategy for a company's software as a service (SaaS) web application running on AWS. For application and security reasons users subscribing to this application are distributed across multiple Application Load Balancers (ALBs) each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances. The application does not require a build stage and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs Auto Scaling groups and EC2 fleets.

Which architecture will meet these requirements with the LEAST amount of configuration?

- A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair.
- B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.
- C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.
- D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and deployment group created for the same ALB-Auto Scaling group pair.

Answer: C

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-groups.html>

NEW QUESTION 106

A company uses AWS Secrets Manager to store a set of sensitive API keys that an AWS Lambda function uses. When the Lambda function is invoked, the Lambda function retrieves the API keys and makes an API call to an external service. The Secrets Manager secret is encrypted with the default AWS Key Management Service (AWS KMS) key.

A DevOps engineer needs to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Update the default KMS key for Secrets Manager to allow only the Lambda function's execution role to decrypt.
- B. Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt
- C. Update Secrets Manager to use the new customer managed key.
- D. Create a KMS customer managed key that trusts Secrets Manager and allows the account's :root principal to decrypt
- E. Update Secrets Manager to use the new customer managed key.
- F. Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource level
- G. Configure the permissions so that the KMS key can encrypt the Secrets Manager secret.
- H. Remove all KMS permissions from the Lambda function's execution role.

Answer: BD

Explanation:

The requirement is to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege, which means granting the minimum permissions necessary to perform a task.

To do this, the DevOps engineer needs to use the following steps:

? Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt. A customer managed key is a symmetric encryption key that is fully managed by the customer. The customer can define the key policy, which specifies who can use and manage the key. By creating a customer managed key, the DevOps engineer can restrict the decryption permission to only the Lambda function's execution role, and prevent other principals from accessing the secret values. The customer managed key also needs to trust Secrets Manager, which means allowing Secrets Manager to use the key to encrypt and decrypt secrets on behalf of the customer.

? Update Secrets Manager to use the new customer managed key. Secrets Manager allows customers to choose which KMS key to use for encrypting each secret. By default, Secrets Manager uses the default KMS key for Secrets Manager, which is a service-managed key that is shared by all customers in the same AWS Region. By updating Secrets Manager to use the new customer managed key, the DevOps engineer can ensure that only the Lambda function's execution role can decrypt the secret values using that key.

? Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource level. The Lambda function's execution role is an IAM role that grants permissions to the Lambda function to access AWS services and resources. The role needs to have KMS permissions to use the customer managed key for decryption. However, to apply the principle of least privilege, the role should have the permissions scoped on the resource level, which means specifying the ARN of the customer managed key as a condition in the IAM policy statement. This way, the role can only use that specific key and not any other KMS keys in the account.

NEW QUESTION 108

A company updated the AWS CloudFormation template for a critical business application. The stack update process failed due to an error in the updated template and AWS CloudFormation automatically began the stack rollback process. Later a DevOps engineer discovered that the application was still unavailable and that the stack was in the UPDATE_ROLLBACK_FAILED state.

Which combination of actions should the DevOps engineer perform so that the stack rollback can complete successfully? (Select TWO.)

- A. Attach the AWS CloudFormation FullAccess IAM policy to the AWS CloudFormation role.
- B. Automatically recover the stack resources by using AWS CloudFormation drift detection.
- C. Issue a ContinueUpdateRollback command from the AWS CloudFormation console or the AWS CLI.
- D. Manually adjust the resources to match the expectations of the stack.
- E. Update the existing AWS CloudFormation stack by using the original template.

Answer: CD

Explanation:

<https://docs.aws.amazon.com/cli/latest/reference/cloudformation/continue-update-rollback.html> For a specified stack that is in the UPDATE_ROLLBACK_FAILED state, continues rolling it back to the UPDATE_ROLLBACK_COMPLETE state. Depending on the cause of the failure, you can manually fix the error and continue the rollback. By continuing the rollback, you can return your stack to a working state (the UPDATE_ROLLBACK_COMPLETE state), and then try to update the stack again.

NEW QUESTION 110

A company needs to implement failover for its application. The application includes an Amazon CloudFront distribution and a public Application Load Balancer (ALB) in an AWS Region. The company has configured the ALB as the default origin for the distribution.

After some recent application outages, the company wants a zero-second RTO. The company deploys the application to a secondary Region in a warm standby configuration. A DevOps engineer needs to automate the failover of the application to the secondary Region so that HTTP GET requests meet the desired RTO. Which solution will meet these requirements?

- A. Create a second CloudFront distribution that has the secondary ALB as the default origin
- B. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both CloudFront distribution
- C. Update the application to use the new record set.
- D. Create a new origin on the distribution for the secondary AL
- E. Create a new origin group
- F. Set the original ALB as the primary origin
- G. Configure the origin group to fail over for HTTP 5xx status code
- H. Update the default behavior to use the origin group.
- I. Create Amazon Route 53 alias records that have a failover policy and Evaluate TargetHealth set to Yes for both ALB
- J. Set the TTL of both records to 0. Update the distribution's origin to use the new record set.
- K. Create a CloudFront function that detects HTTP 5xx status code
- L. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status code
- M. Update the distribution's default behavior to send origin responses to the function.

Answer: B

Explanation:

The best solution to implement failover for the application is to use CloudFront origin groups. Origin groups allow CloudFront to automatically switch to a secondary origin when the primary origin is unavailable or returns specific HTTP status codes that indicate a failure¹. This way, CloudFront can serve the requests from the secondary ALB in the secondary Region without any delay or redirection. To set up origin groups, the DevOps engineer needs to create a new origin on the distribution for the secondary ALB, create a new origin group with the original ALB as the primary origin and the secondary ALB as the secondary origin, and configure the origin group to fail over for HTTP 5xx status

codes. Then, the DevOps engineer needs to update the default behavior to use the origin group instead of the single origin².

The other options are not as effective or efficient as the solution in option B. Option A is not suitable because creating a second CloudFront distribution will increase the complexity and cost of the application. Moreover, using Route 53 alias records with a failover policy will introduce some delay in detecting and switching to the secondary CloudFront distribution, which may not meet the zero-second RTO requirement. Option C is not feasible because CloudFront does not support using Route 53 alias records as origins³. Option D is not advisable because using a CloudFront function to redirect the requests to the secondary ALB will add an extra round-trip and latency to the failover process, which may also not meet the zero-second RTO requirement.

References:

? 1: Optimizing high availability with CloudFront origin failover - Amazon CloudFront

? 2: Creating an origin group - Amazon CloudFront

? 3: Values That You Specify When You Create or Update a Web Distribution - Amazon CloudFront

NEW QUESTION 113

A company's application uses a fleet of Amazon EC2 On-Demand Instances to analyze and process data. The EC2 instances are in an Auto Scaling group. The Auto Scaling group is a target group for an Application Load Balancer (ALB). The application analyzes critical data that cannot tolerate interruption. The application also analyzes noncritical data that can withstand interruption.

The critical data analysis requires quick scalability in response to real-time application demand. The noncritical data analysis involves memory consumption. A DevOps engineer must implement a solution that reduces scale-out latency for the critical data. The solution also must process the noncritical data.

Which combination of steps will meet these requirements? (Select TWO.)

- A. For the critical data, modify the existing Auto Scaling group
- B. Create a warm pool instance in the stopped state
- C. Define the warm pool size
- D. Create a new version of the launch template that has detailed monitoring enabled
- E. Use Spot Instances.
- F. For the critical data, modify the existing Auto Scaling group
- G. Create a warm pool instance in the stopped state
- H. Define the warm pool size
- I. Create a new version of the launch template that has detailed monitoring enabled
- J. Use On-Demand Instances.
- K. For the critical data
- L. modify the existing Auto Scaling group
- M. Create a lifecycle hook to ensure that bootstrap scripts are completed successfully
- N. Ensure that the application on the instances is ready to accept traffic before the instances are registered
- O. Create a new version of the launch template that has detailed monitoring enabled.
- P. For the noncritical data, create a second Auto Scaling group that uses a launch template
- Q. Configure the launch template to install the unified Amazon CloudWatch agent and to configure the CloudWatch agent with a custom memory utilization metric
- R. Use Spot Instance
- S. Add the new Auto Scaling group as the target group for the ALB
- T. Modify the application to use two target groups for critical data and noncritical data.
- . For the noncritical data, create a second Auto Scaling group
- . Choose the predefined memory utilization metric type for the target tracking scaling policy
- . Use Spot Instance
- . Add the new Auto Scaling group as the target group for the ALB
- . Modify the application to use two target groups for critical data and noncritical data.

Answer: BD

Explanation:

? For the critical data, using a warm pool¹ can reduce the scale-out latency by having pre-initialized EC2 instances ready to serve the application traffic. Using On-Demand Instances can ensure that the instances are always available and not interrupted by Spot interruptions².

? For the noncritical data, using a second Auto Scaling group with Spot Instances can reduce the cost and leverage the unused capacity of EC2³. Using a launch template with the CloudWatch agent⁴ can enable the collection of memory utilization metrics, which can be used to scale the group based on the memory demand. Adding the second group as a target group for the ALB and modifying the application to use two target groups can enable routing the traffic based on the data type.

References: 1: Warm pools for Amazon EC2 Auto Scaling 2: Amazon EC2 On-Demand Capacity Reservations 3: Amazon EC2 Spot Instances 4: Metrics collected by the CloudWatch agent

NEW QUESTION 116

A company uses AWS Directory Service for Microsoft Active Directory as its identity provider (IdP). The company requires all infrastructure to be defined and deployed by AWS CloudFormation.

A DevOps engineer needs to create a fleet of Windows-based Amazon EC2 instances to host an application. The DevOps engineer has created a CloudFormation template that contains an EC2 launch template, IAM role, EC2 security group, and EC2 Auto Scaling group. The DevOps engineer must implement a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory.

Which solution will meet these requirements with the MOST operational efficiency?

- A. In the CloudFormation template, create an AWS::SSM::Document resource that joins the EC2 instance to the AWS Managed Microsoft AD domain by using the parameters for the existing director
- B. Update the launch template to include the SSMAssociation property to use the new SSM document
- C. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.
- D. In the CloudFormation template, update the launch template to include specific tags that propagate on launch
- E. Create an AWS::SSM::Association resource to associate the AWS- JoinDirectoryServiceDomain Automation runbook with the EC2 instances that have the specified tag
- F. Define the required parameters to join the AWS Managed Microsoft AD director
- G. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.
- H. Store the existing AWS Managed Microsoft AD domain connection details in AWS Secrets Manager

- I. In the CloudFormation template, create an AWS::SSM::Association resource to associate the AWS-CreateManagedWindowsInstanceWithApproval Automation runbook with the EC2 Auto Scaling group
- J. Pass the ARNs for the parameters from Secrets Manager to join the domain
- K. Attach the AmazonSSMDirectoryServiceAccess and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.
- L. Store the existing AWS Managed Microsoft AD domain administrator credentials in AWS Secrets Manager
- M. In the CloudFormation template, update the EC2 launch template to include user data
- N. Configure the user data to pull the administrator credentials from Secrets Manager and to join the AWS Managed Microsoft AD domain
- O. Attach the AmazonSSMManagedInstanceCore and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.

Answer: B

Explanation:

To meet the requirements, the DevOps engineer needs to create a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory with the most operational efficiency. The DevOps engineer can use AWS Systems Manager Automation to automate the domain join process using an existing runbook called AWS- JoinDirectoryServiceDomain. This runbook can join Windows instances to an AWS Managed Microsoft AD or Simple AD directory by using PowerShell commands. The DevOps engineer can create an AWS::SSM::Association resource in the CloudFormation template to associate the runbook with the EC2 instances that have specific tags. The tags can be defined in the launch template and propagated on launch to the EC2 instances. The DevOps engineer can also define the required parameters for the runbook, such as the directory ID, directory name, and organizational unit. The DevOps engineer can attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use. These policies grant the necessary permissions for Systems Manager and Directory Service operations.

NEW QUESTION 117

A company is implementing an Amazon Elastic Container Service (Amazon ECS) cluster to run its workload. The company architecture will run multiple ECS services on the cluster. The architecture includes an Application Load Balancer on the front end and uses multiple target groups to route traffic.

A DevOps engineer must collect application and access logs. The DevOps engineer then needs to send the logs to an Amazon S3 bucket for near-real-time analysis.

Which combination of steps must the DevOps engineer take to meet these requirements? (Choose three.)

- A. Download the Amazon CloudWatch Logs container instance from AWS
- B. Configure this instance as a task
- C. Update the application service definitions to include the logging task.
- D. Install the Amazon CloudWatch Logs agent on the ECS instance
- E. Change the logging driver in the ECS task definition to awslogs.
- F. Use Amazon EventBridge to schedule an AWS Lambda function that will run every 60 seconds and will run the Amazon CloudWatch Logs create-export-task command
- G. Then point the output to the logging S3 bucket.
- H. Activate access logging on the ALB
- I. Then point the ALB directly to the logging S3 bucket.
- J. Activate access logging on the target groups that the ECS services use
- K. Then send the logs directly to the logging S3 bucket.
- L. Create an Amazon Kinesis Data Firehose delivery stream that has a destination of the logging S3 bucket
- M. Then create an Amazon CloudWatch Logs subscription filter for Kinesis Data Firehose.

Answer: BDF

Explanation:

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-logging-monitoring.html>

NEW QUESTION 122

A DevOps engineer manages a large commercial website that runs on Amazon EC2. The website uses Amazon Kinesis Data Streams to collect and process web logs. The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2.

Sudden increases of data cause the Kinesis consumer application to fall behind and the Kinesis data streams drop records before the records can be processed.

The DevOps engineer must implement a solution to improve stream handling.

Which solution meets these requirements with the MOST operational efficiency?

- A. Modify the Kinesis consumer application to store the logs durably in Amazon S3. Use Amazon EMR to process the data directly on Amazon S3 to derive customer insights. Store the results in Amazon S3.
- B. Horizontally scale the Kinesis consumer application by adding more EC2 instances based on the Amazon CloudWatch GetRecords.IteratorAge.Milliseconds metric. Increase the retention period of the Kinesis data streams.
- C. Convert the Kinesis consumer application to run as an AWS Lambda function
- D. Configure the Kinesis data streams as the event source for the Lambda function to process the data streams
- E. Increase the number of shards in the Kinesis data streams to increase the overall throughput so that the consumer application processes the data faster.

Answer: B

Explanation:

<https://docs.aws.amazon.com/streams/latest/dev/monitoring-with-cloudwatch.html>

GetRecords.IteratorAge.Milliseconds - The age of the last record in all GetRecords calls made against a Kinesis stream, measured over the specified time period. Age is the difference between the current time and when the last record of the GetRecords call was written to the stream. The Minimum and Maximum statistics can be used to track the progress of Kinesis consumer applications. A value of zero indicates that the records being read are completely caught up.

NEW QUESTION 124

A company plans to use Amazon CloudWatch to monitor its Amazon EC2 instances. The company needs to stop EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. The company must evaluate the metric every hour. The EC2 instances must continue to run if there is missing data for the NetworkPacketsIn metric during the evaluation period.

A DevOps engineer creates a CloudWatch alarm for the NetworkPacketsIn metric. The DevOps engineer configures a threshold value of 5 and an evaluation period of 1 hour.

Which set of additional actions should the DevOps engineer take to meet these requirements?

- A. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as breaching the threshold

- B. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.
- C. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as not breaching the threshold.
- D. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
- E. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as breaching the threshold.
- F. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
- G. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as not breaching the threshold.
- H. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.

Answer: B

Explanation:

To meet the requirements, the DevOps engineer needs to configure the CloudWatch alarm to stop the EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. This means that the alarm should trigger when 3 out of 12 datapoints are below the threshold of 5. The alarm should also treat missing data as not breaching the threshold, so that the EC2 instances continue to run if there is no data for the metric during the evaluation period. The DevOps engineer can add an EC2 action to stop the instance when the alarm enters the ALARM state, which is a built-in action type for CloudWatch alarms.

NEW QUESTION 129

A DevOps engineer is setting up a container-based architecture. The engineer has decided to use AWS CloudFormation to automatically provision an Amazon ECS cluster and an Amazon EC2 Auto Scaling group to launch the EC2 container instances. After successfully creating the CloudFormation stack, the engineer noticed that, even though the ECS cluster and the EC2 instances were created successfully and the stack finished the creation, the EC2 instances were associating with a different cluster.

How should the DevOps engineer update the CloudFormation template to resolve this issue?

- A. Reference the EC2 instances in the AWS: ECS: Cluster resource and reference the ECS cluster in the AWS: ECS: Service resource.
- B. Reference the ECS cluster in the AWS: AutoScaling: LaunchConfiguration resource of the UserData property.
- C. Reference the ECS cluster in the AWS:EC2: Instance resource of the UserData property.
- D. Reference the ECS cluster in the AWS: CloudFormation: CustomResource resource to trigger an AWS Lambda function that registers the EC2 instances with the appropriate ECS cluster.

Answer: B

Explanation:

The UserData property of the AWS: AutoScaling: LaunchConfiguration resource can be used to specify a script that runs when the EC2 instances are launched. This script can include the ECS cluster name as an environment variable for the ECS agent running on the EC2 instances. This way, the EC2 instances will register with the correct ECS cluster. Option A is incorrect because the AWS: ECS: Cluster resource does not have a property to reference the EC2 instances. Option C is incorrect because the EC2 instances are launched by the Auto Scaling group, not by the AWS: EC2: Instance resource. Option D is incorrect because using a custom resource and a Lambda function is unnecessary and overly complex for this scenario. References: AWS::AutoScaling::LaunchConfiguration, Amazon ECS Container Agent Configuration

NEW QUESTION 132

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DOP-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DOP-C02 Product From:

<https://www.2passeasy.com/dumps/DOP-C02/>

Money Back Guarantee

DOP-C02 Practice Exam Features:

- * DOP-C02 Questions and Answers Updated Frequently
- * DOP-C02 Practice Questions Verified by Expert Senior Certified Staff
- * DOP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DOP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year