Amazon-Web-Services

Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional

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 maintenanc
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 windo
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 maintenanc
F. Target a newly created AWS Lambda function thatregisters 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 maintenanc
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 to 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 AmazonGuardDuty administrator accoun
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 Cloud Formation stack set that accepts the GuardDuty invitation and creates an Amazon EventBridge rule Configure the rule with an event pattern to matc
D. GuardDuty events and to forward matching events to the SNS topi
E. Configure the Cloud Formation stack set to deploy into all AWS accounts in the organization.
F. In the organization's management accoun
G. create an AWS CloudTrail organization trail Activate the organization trail in all AWS accounts in the organizatio
H. Create an SCP that enables VPC Flow Logs in each account in the organizatio
I. Configure AWS Security Hub for the organization Create an Amazon EventBridge rule with an even 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 trai
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 organizatio
L. Configure AWS Security Hub for the organizatio
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 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 4
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 5
A company uses AWS Key Management Service (AWS KMS) keys and manual key rotation to meet regulatory compliance requirements. The security team wants to be notified when any keys have not been rotated after 90 days.
Which solution will accomplish this?

A. Configure AWS KMS to publish to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.
B. Configure an Amazon EventBridge event to launch an AWS Lambda function to call the AWS Trusted Advisor API and publish to an Amazon Simple Notification Service (Amazon SNS) topic.
C. Develop an AWS Config custom rule that publishes to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.
D. Configure AWS Security Hub to publish to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.

Answer: C

Explanation:
https://aws.amazon.com/blogs/security/how-to-use-aws-config-to-determine-compliance-of-aws-kms-key-policies-to-your-specifications/

NEW QUESTION 6
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 Regio

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 Regio
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 Regio
H. Retrieve the data from a DynamoDB global tabl
I. Deploy a Lambda function to check the North America API health every 5 minute
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 routin
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 7
A company has a new AWS account that teams will use to deploy various applications. The teams will create many Amazon S3 buckets for application- specific purposes and to store AWS CloudTrail logs. The company has enabled Amazon Macie for the account.
A DevOps engineer needs to optimize the Macie costs for the account without compromising the account's functionality.
Which solutions will meet these requirements? (Select TWO.)

A. Exclude S3 buckets that contain CloudTrail logs from automated discovery.
B. Exclude S3 buckets that have public read access from automated discovery.
C. Configure scheduled daily discovery jobs for all S3 buckets in the account.
D. Configure discovery jobs to include S3 objects based on the last modified criterion.
E. Configure discovery jobs to include S3 objects that are tagged as production only.

Answer: AD

Explanation:
To optimize the Macie costs for the account without compromising the account's functionality, the DevOps engineer needs to exclude S3 buckets that do not contain sensitive data from automated discovery. S3 buckets that contain CloudTrail logs are unlikely to have sensitive data, and Macie charges for scanning and monitoring data in S3 buckets. Therefore, excluding S3 buckets that contain CloudTrail logs from automated discovery can reduce Macie costs. Similarly, configuring discovery jobs to include S3 objects based on the last modified criterion can also reduce Macie costs, as it will only scan and monitor new or updated objects, rather than all objects in the bucket.

NEW QUESTION 8
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 grou
B. Configure a CloudWatch Logs metric filter that increments a metric for each API operation nam
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 grou
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 grou
H. Modify the Lambda function to respond to the ALB with the API operation name, response code, and version number as response metadat
I. Configure a CloudWatch Logs metric filter that increments a metric for each API operation nam
J. Specify response code and application version as dimensions for the metric.
K. Configure AWS X-Ray integration on the Lambda functio
L. Modify the Lambda function to create an X-Ray subsegment with the API operation name, response code, and version numbe
M. Configure X-Ray insights to extract an aggregated metric for each API operation name and to publish the metric to Amazon CloudWatc
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 9
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 rol
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 I

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 accoun
G. Attach a trust policy that allows logs.amazonaws.com to perform the iam:CreateSink action if the aws:PrincipalOrgld 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 sin
J. Allow logs to be shared with the monitoring accoun
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 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 installe
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 installe
D. Create an AppSpec file that contains application deployment scripts and grants access to CodeDeploy.
E. Create an application in CodeDeplo
F. Configure an in-place deployment typ
G. Specify the Auto Scaling group as the deployment targe
H. Add a step to the CodePipeline pipeline to use EC2 Image Builder to create a new AM
I. Configure CodeDeploy to deploy the newly created AMI.
J. Create an application in CodeDeplo
K. Configure an in-place deployment typ
L. Specify the Auto Scaling group as the deployment targe
M. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.
N. Create an application in CodeDeplo
O. Configure an in-place deployment typ
P. Specify the EC2 instances that are launched from the common AMI as the deployment targe
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 10
A large enterprise is deploying a web application on AWS. The application runs on Amazon
EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS for Oracle DB instance and Amazon DynamoDB. There are separate environments tor development testing and production.
What is the MOST secure and flexible way to obtain password credentials during deployment?

A. Retrieve an access key from an AWS Systems Manager securestring parameter to access AWS service
B. Retrieve the database credentials from a Systems Manager SecureString parameter.
C. Launch the EC2 instances with an EC2 1AM role to access AWS services Retrieve the database credentials from AWS Secrets Manager.
D. Retrieve an access key from an AWS Systems Manager plaintext parameter to access AWS service
E. Retrieve the database credentials from a Systems Manager SecureString parameter.
F. Launch the EC2 instances with an EC2 1AM role to access AWS services Store the database passwords in an encrypted config file with the application artifacts.

Answer: B

Explanation:
AWS Secrets Manager is a secrets management service that helps you protect access to your applications, services, and IT resources. This service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Using Secrets Manager, you can secure and manage secrets used to access resources in the AWS Cloud, on third-party services, and on-premises. SSM parameter store and AWS Secret manager are both a secure option. However, Secrets manager is more flexible and has more options like password generation. Reference:
https://www.1strategy.com/blog/2019/02/28/aws-parameter-store-vs-aws- secrets-manager/

NEW QUESTION 12
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 Semce (Amazon SNS) topic and subscribe 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 interva
C. Configure the canary to use the SNS topic when the applications return errors.
D. Create an Amazon CloudWatch alarm that uses the AWS/AppljcabonELB 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 13
A company manages multiple AWS accounts by using AWS Organizations with OUS for the different business divisions, The company is updating their corporate network to use new IP address ranges. The company has 10 Amazon S3 buckets in different AWS accounts. The S3 buckets store reports for the different divisions. The S3 bucket configurations allow only private corporate network IP addresses to access the S3 buckets.
A DevOps engineer needs to change the range of IP addresses that have permission to access the contents of the S3 buckets The DevOps engineer also needs to revoke the permissions of two OUS in the company
Which solution will meet these requirements?

A. Create a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that demes access to the old range of IP addresses for all the S3 bucket
B. Set a permissions boundary for the OrganzauonAccountAccessRole role In the two OUS to deny access to the S3 buckets.
C. Create a new SCP that has a statement that allows only the new range of IP addresses to access the S3 bucket
D. Create another SCP that denies access to the S3 bucket
E. Attach the second SCP to the two OUS
F. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket
G. Create a new SCP that denies access to the S3 bucket
H. Attach the SCP to the two OUs.
I. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket
J. Set a permissions boundary for the OrganizationAccountAccessRole role in the two OUS to deny access to the S3 buckets.

Answer: C

Explanation:
The correct answer is C.
A comprehensive and detailed explanation is:
? Option A is incorrect because creating a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that denies access to the old range of IP addresses for all the S3 buckets, is not a valid solution. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. SCPs can only control the actions that can be performed by the principals in the organization, not the access to specific resources. Moreover, setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets.
? Option B is incorrect because creating a new SCP that has a statement that allows
only the new range of IP addresses to access the S3 buckets is not a valid solution, for the same reason as option A. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. Creating another SCP that denies access to the S3 buckets and attaching it to the two OUs is also not a valid solution, as SCPs cannot specify the S3 buckets as resources either.
? Option C is correct because it meets both requirements of changing the range of IP addresses that have permission to access the contents of the S3 buckets and revoking the permissions of two OUs in the company. On all the S3 buckets, configuring resource-based policies that allow only the new range of IP addresses to access the S3 buckets is a valid way to update the IP address ranges, as resource-based policies can specify both resources and conditions. Creating a new SCP that denies access to the S3 buckets and attaching it to the two OUs is also a valid way to revoke the permissions of those OUs, as SCPs can deny actions such as s3:PutObject or s3:GetObject on any resource.
? Option D is incorrect because setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets. A permissions boundary is a policy that defines the maximum permissions that an IAM entity can have. However, it does not revoke any existing permissions that are granted by other policies.
References:

? AWS Organizations
? S3 Bucket Policies
? Service Control Policies
? Permissions Boundaries

NEW QUESTION 16
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 accoun
B. Attach the AdministratorAccess policy to the rol
C. Modify the trust relationship to allow the sts:AssumeRole action from the workload accounts.
D. Create a SysAdmin role in each workload accoun
E. Attach the AdministratorAccess policy to the rol
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 accoun
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 accoun
L. Add all operations team members to the group.
M. Create an Amazon Cognito user pool in the operations accoun
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 18
A healthcare services company is concerned about the growing costs of software licensing for an application for monitoring patient wellness. The company wants to create an audit process to ensure that the application is running exclusively on Amazon EC2 Dedicated Hosts. A DevOps engineer must create a workflow to audit the application to ensure compliance.
What steps should the engineer take to meet this requirement with the LEAST administrative overhead?

A. Use AWS Systems Manager Configuration Complianc
B. Use calls to the put- compliance-items API action to scan and build a database of noncompliant EC2 instances based on their host placement configuratio
C. Use an Amazon DynamoDB table to store these instance IDs for fast acces
D. Generate a report through Systems Manager by calling the list-compliance-summaries API action.
E. Use custom Java code running on an EC2 instanc
F. Set up EC2 Auto Scaling for the instance depending on the number of instances to be checke
G. Send the list of noncompliant EC2 instance IDs to an Amazon SQS queu
H. Set up another worker instance to process instance IDs from the SQS queue and write them to Amazon DynamoD
I. Use an AWS Lambda function to terminate noncompliant instance IDs obtained from the queue, and send them to an Amazon SNS email topic for distribution.
J. Use AWS Confi
K. Identify all EC2 instances to be audited by enabling Config Recording on all Amazon EC2 resources for the regio
L. Create a custom AWS Config rule that triggers an AWS Lambda function by using the "config-rule-change-triggered" blueprint.Modify the LambdaevaluateCompliance () function to verify host placement to return a NON_COMPLIANT result if the instance is not running on an EC2 Dedicated Hos
M. Use the AWS Config report to address noncompliant instances.
N. Use AWS CloudTrai
O. Identify all EC2 instances to be audited by analyzing all calls to the EC2 RunCommand API actio
P. Invoke a AWS Lambda function that analyzes the host placement of the instanc
Q. Store the EC2 instance ID of noncompliant resources in an Amazon RDS for MySQL DB instanc
R. Generate a report by querying the RDS instance and exporting the query results to a CSV text file.

Answer: C

Explanation:
The correct answer is C. Using AWS Config to identify and audit all EC2 instances based on their host placement configuration is the most efficient and scalable solution to ensure compliance with the software licensing requirement. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. By creating a custom AWS Config rule that triggers a Lambda function to verify host placement, the DevOps engineer can automate the process of checking whether the instances are running on EC2 Dedicated Hosts or not. The Lambda function can return a NON_COMPLIANT result if the instance is not running on an EC2 Dedicated Host, and the AWS Config report can provide a summary of the compliance status of the instances. This solution requires the least administrative overhead compared to the other options.
Option A is incorrect because using AWS Systems Manager Configuration Compliance to scan and build a database of noncompliant EC2 instances based on their host placement configuration is a more complex and costly solution than using AWS Config. AWS Systems Manager Configuration Compliance is a feature of AWS Systems Manager that enables you to scan your managed instances for patch compliance and configuration inconsistencies. To use this feature, the DevOps engineer would need to install the Systems Manager Agent on each EC2 instance, create a State Manager association to run the put-compliance-items API action periodically, and use a DynamoDB table to store the instance IDs of noncompliant resources. This solution would also require more API calls and storage costs than using AWS Config.
Option B is incorrect because using custom Java code running on an EC2 instance to check and terminate noncompliant EC2 instances is a more cumbersome and error-prone solution than using AWS Config. This solution would require the DevOps engineer to write and maintain the Java code, set up EC2 Auto Scaling for the instance, use an SQS queue and another worker instance to process the instance IDs, use a Lambda function and an SNS topic to terminate and notify the noncompliant instances, and handle any potential failures or exceptions in the workflow. This solution would also incur more compute,
storage, and messaging costs than using AWS Config.
Option D is incorrect because using AWS CloudTrail to identify and audit EC2 instances by analyzing the EC2 RunCommand API action is a less reliable and accurate solution than using AWS Config. AWS CloudTrail is a service that enables you to monitor and log the API activity in your AWS account. The EC2 RunCommand API action is used to execute commands on one or more EC2 instances. However, this API action does not necessarily indicate the host placement of the instance, and it may not capture all the instances that are running on EC2 Dedicated Hosts or not. Therefore, option D would not provide a comprehensive

and consistent audit of the EC2 instances.

NEW QUESTION 20
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 accoun
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 accoun
G. Configure Amazon Inspector to activate deep inspection for EC2 instance
H. Create an Amazon EventBridge rule for an Inspector2 findin
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 even
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 21
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 topi
B. Create an AWS CloudTrail subscription for the SNS topi
C. Configure a subscription filter for security group modification events.
D. Create an Amazon EventBridge scheduled rule to invoke the Lambda functio
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 sourc
G. Define the rule’s event pattern to match EC2 security group creation and modification event
H. Configure the rule to invoke the Lambda function.
I. Create an Amazon EventBridge custom event bus that subscribes to events from all AWS service
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 22
A company has multiple member accounts that are part of an organization in AWS Organizations. The security team needs to review every Amazon EC2 security group and their inbound and outbound rules. The security team wants to programmatically retrieve this information from the member accounts using an AWS Lambda function in the management account of the organization.
Which combination of access changes will meet these requirements? (Choose three.)

A. Create a trust relationship that allows users in the member accounts to assume the management account IAM role.
B. Create a trust relationship that allows users in the management account to assume the IAM roles of the member accounts.
C. Create an IAM role in each member account that has access to the AmazonEC2ReadOnlyAccess managed policy.
D. Create an I AM role in each member account to allow the sts:AssumeRole action against the management account IAM role's ARN.
E. Create an I AM role in the management account that allows the sts:AssumeRole action against the member account IAM role's ARN.
F. Create an IAM role in the management account that has access to the AmazonEC2ReadOnlyAccess managed policy.

Answer: BCE

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/lambda-function-assume-iam-role/ https://kreuzwerker.de/post/aws-multi-account-setups-reloaded

NEW QUESTION 26
A company sells products through an ecommerce web application The company wants a dashboard that shows a pie chart of product transaction details. The company wants to integrate the dashboard With the company’s existing Amazon CloudWatch dashboards
Which solution Will meet these requirements With the MOST operational efficiency?

A. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transactio
B. Use CloudWatch Logs Insights to query the log group and to visualize the results in a pie chart format Attach the results to the desired CloudWatch dashboard.
C. Update the ecommerce application to emit a JSON object to an Amazon S3 bucket for each processed transactio
D. Use Amazon Athena to query the S3 bucket and to visualize the results In a Pie chart forma
E. Export the results from Athena Attach the results to the desired CloudWatch dashboard
F. Update the ecommerce application to use AWS X-Ray for instrumentatio
G. Create a new X-Ray subsegment Add an annotation for each processed transactio
H. Use X-Ray traces to query the data and to visualize the results in a pie chart format Attach the results to the desired CloudWatch dashboard
I. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction_ Create an AWS Lambda function to aggregate and write the results to Amazon DynamoD
J. Create a Lambda subscription filter for the log fil
K. Attach the results to the desired CloudWatch dashboard.

Answer: A

Explanation:
The correct answer is A.
A comprehensive and detailed explanation is:
? Option A is correct because it meets the requirements with the most operational efficiency. Updating the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction is a simple and cost- effective way to collect the data needed for the dashboard. Using CloudWatch Logs Insights to query the log group and to visualize the results in a pie chart format is also a convenient and integrated solution that leverages the existing CloudWatch dashboards. Attaching the results to the desired CloudWatch dashboard is straightforward and does not require any additional steps or services.
? Option B is incorrect because it introduces unnecessary complexity and cost.
Updating the ecommerce application to emit a JSON object to an Amazon S3 bucket for each processed transaction is a valid way to store the data, but it requires creating and managing an S3 bucket and its permissions. Using Amazon Athena to query the S3 bucket and to visualize the results in a pie chart format is also a valid way to analyze the data, but it incurs charges based on the amount of
data scanned by each query. Exporting the results from Athena and attaching them to the desired CloudWatch dashboard is also an extra step that adds more overhead and latency.
? Option C is incorrect because it uses AWS X-Ray for an inappropriate purpose.
Updating the ecommerce application to use AWS X-Ray for instrumentation is a good practice for monitoring and tracing distributed applications, but it is not designed for aggregating product transaction details. Creating a new X-Ray subsegment and adding an annotation for each processed transaction is possible, but it would clutter the X-Ray service map and make it harder to debug performance issues. Using X-Ray traces to query the data and to visualize the results in a pie chart format is also possible, but it would require custom code and logic that are not supported by X-Ray natively. Attaching the results to the desired CloudWatch dashboard is also not supported by X-Ray directly, and would require additional steps or services.
? Option D is incorrect because it introduces unnecessary complexity and cost.
Updating the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction is a simple and cost-effective way to collect the data needed for the dashboard, as in option A. However, creating an AWS Lambda function to aggregate and write the results to Amazon DynamoDB is redundant, as CloudWatch Logs Insights can already perform aggregation queries on log data. Creating a Lambda subscription filter for the log file is also redundant, as CloudWatch Logs Insights can already access log data directly. Attaching the results to the desired CloudWatch dashboard would also require additional steps or services, as DynamoDB does not support native integration with CloudWatch dashboards.
References:
? CloudWatch Logs Insights
? Amazon Athena
? AWS X-Ray
? AWS Lambda
? Amazon DynamoDB

NEW QUESTION 29
A DevOps engineer is creating an AWS CloudFormation template to deploy a web service. The web service will run on Amazon EC2 instances in a private subnet behind an Application Load Balancer (ALB). The DevOps engineer must ensure that the service can accept requests from clients that have IPv6 addresses.
What should the DevOps engineer do with the CloudFormation template so that IPv6 clients can access the web service?

A. Add an IPv6 CIDR block to the VPC and the private subnet for the EC2 instance
B. Create route table entries for the IPv6 network, use EC2 instance types that support IPv6, and assign IPv6 addresses to each EC2 instance.
C. Assign each EC2 instance an IPv6 Elastic IP addres
D. Create a target group, and add the EC2 instances as target
E. Create a listener on port 443 of the ALB, and associate the target group with the ALB.
F. Replace the ALB with a Network Load Balancer (NLB). Add an IPv6 CIDR block to the VPC and subnets for the NLB, and assign the NLB an IPv6 Elastic IP address.
G. Add an IPv6 CIDR block to the VPC and subnets for the AL
H. Create a listener on port 443. and specify the dualstack IP address type on the AL
I. Create a target group, and add the EC2 instances as target
J. Associate the target group with the ALB.

Answer: D

Explanation:
it involves adding an IPv6 CIDR block to the VPC and subnets for the ALB and specifying the dualstack IP address type on the ALB listener. This allows the ALB to listen on both IPv4 and IPv6 addresses, and forward requests to the EC2 instances that are added as targets to the target group associated with the ALB.

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 numb 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 togs The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2 Sudden increases of data cause the Kinesis consumer application to (all 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'' er of pull requests for certain files.
How can the company meet these requirements with the LEAST amount of effort?

A. Activate S3 server access loggin

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 loggin
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 35
An Amazon EC2 instance is running in a VPC and needs to download an object from a restricted Amazon S3 bucket. When the DevOps engineer tries to download the object, an AccessDenied error is received,
What are the possible causes tor this error? (Select TWO,)

A. The 53 bucket default encryption is enabled.
B. There is an error in the S3 bucket policy.
C. The object has been moved to S3 Glacier.
D. There is an error in the IAM role configuration.
E. S3 Versioning is enabled.

Answer: BD

Explanation:
These are the possible causes for the AccessDenied error because they affect the permissions to access the S3 object from the EC2 instance. An S3 bucket policy is a resource-based policy that defines who can access the bucket and its objects, and what actions they can perform. An IAM role is an identity that can be assumed by an EC2 instance to grant it permissions to access AWS services and resources. If there is an error in the S3 bucket policy or the IAM role configuration, such as a missing or incorrect statement, condition, or principal, then the EC2 instance may not have the necessary permissions to download the object from the S3 bucket . https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies.html
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

NEW QUESTION 37
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 38
A DevOps team uses AWS CodePipeline, AWS CodeBuild, and AWS CodeDeploy to deploy an application. The application is a REST API that uses AWS Lambda functions and Amazon API Gateway Recent deployments have introduced errors that have affected many customers.
The DevOps team needs a solution that reverts to the most recent stable version of the application when an error is detected. The solution must affect the fewest customers possible.
Which solution Will meet these requirements With the MOST operational efficiency?

A. Set the deployment configuration in CodeDepIoy to LambdaAllAtOnce Configure automatic rollbacks on the deployment group Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway Configure the deployment group to roll back when the number of alarms meets the alarm threshold
B. Set the deployment configuration in CodeDeploy to LambdaCanary10Percent10Minute
C. Configure automatic rollbacks on the deployment group Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway Configure the deployment group to roll back when the number of alarms meets the alarm threshold
D. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce Configure manual rollbacks on the deployment grou
E. Create an Amazon Simple Notification Service (Amazon SNS) topc to send notifications every time a deploymrent fad
F. Configure the SNS topc to Invoke a new Lambda function that stops the current deployment and starts the most recent successful deployment
G. Set the deployment configuration in CodeDeploy to LambdaCanaryIOPercentIOMinutes Configure manual rollbacks on the deployment group Create a metric filter on an Amazon CloudWatch log group for API Gateway to monitor HTTP Bad Gateway error
H. Configure the metric filter to Invoke a new Lambda function that stops the current eployment and starts the most recent successful deployment

Answer: B

Explanation:
? Option A is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring automatic rollbacks on the deployment group is not operationally efficient, as it requires manual intervention to fix the errors and redeploy the application.
? Option B is correct because setting the deployment configuration to LambdaCanary10Percent10Minutes means that the new version of the application will be deployed to 10 percent of the Lambda functions first, and then to the remaining 90 percent after 10 minutes. This minimizes the impact of errors on customers, as only 10 percent of them will be affected by a faulty deployment. Configuring automatic rollbacks on the deployment group also meets the requirement of reverting to the most recent stable version of the application when an error is detected. Creating a CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway is a valid way to monitor the health of the application and trigger a rollback if needed.
? Option C is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating an SNS topic to send notifications every time a deployment fails is not sufficient to detect errors in the application, as it does not monitor the API Gateway responses.
? Option D is incorrect because configuring manual rollbacks on the deployment
group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating a metric filter on a CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors is a valid way to monitor the health of the application, but invoking a new Lambda function to perform a rollback is unnecessary and complex, as CodeDeploy already provides automatic rollback functionality.
References:
? AWS CodeDeploy Deployment Configurations
? [AWS CodeDeploy Rollbacks]
? Amazon CloudWatch Alarms

NEW QUESTION 42
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 47
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 Hu
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 traffi
G. Configure automatic remediation to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
H. Enable Amazon Inspecto
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 50
A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.
Which solution will accomplish this with the LEAST amount of development effort?

A. Create an Amazon CloudWatch Event rule that runs periodically and targets an AWS Lambda functio
B. Within the Lambda function, evaluate the current state of the AWS environment and compare deployed resource values to resource limits on the accoun
C. Notify the Senior Manager if the account is approaching a service limit.
D. Deploy an AWS Lambda function that refreshes AWS Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodicall
E. Create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda functio
F. In the target Lambda function, notify the Senior Manager.
G. Deploy an AWS Lambda function that refreshes AWS Personal Health Dashboard checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodicall
H. Create another CloudWatch Events rule with an event pattern matching Personal Health Dashboard events and a target Lambda functio
I. In the target Lambda function, notify the Senior Manager.
J. Add an AWS Config custom rule that runs periodically, checks the AWS service limit status, and streams notifications to an Amazon SNS topi
K. Deploy an AWS Lambda function that notifies the Senior Manager, and subscribe the Lambda function to the SNS topic.

Answer: B

Explanation:
To meet the requirements, the company needs to create a solution that alerts the Senior Manager when the creation of resources approaches the service limits for the account with the least amount of development effort. The company can use AWS Trusted Advisor, which is a service that provides best practice recommendations for cost optimization, performance, security, and service limits. The company can deploy an AWS Lambda function that refreshes Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically. This will ensure that Trusted Advisor checks are up to date and reflect the current state of the account. The company can then create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda function. The event pattern can filter for events related to service limit checks and their status. The target Lambda function can notify the Senior Manager via a third-party API call if the event indicates that the account is approaching or exceeding a service limit.

NEW QUESTION 52
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 R TO.
Which solution will meet these requirements?

A. Create a second CloudFront distribution that has the secondary ALB as the default origi
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 grou
F. Set the original ALB as the primary origi
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 Target Health set to Yes for both ALB

J. Set the TTL of both records to
K. Update the distribution's origin to use the new record set.
L. Create a CloudFront function that detects HTTP 5xx status code
M. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status code
N. Update the distribution's default behavior to send origin responses to the function.

Answer: B

Explanation:
To implement failover for the application to the secondary Region so that HTTP GET requests meet the desired RTO, the DevOps engineer should use the following solution:
? Create a new origin on the distribution for the secondary ALB. A CloudFront origin
is the source of the content that CloudFront delivers to viewers. By creating a new origin for the secondary ALB, the DevOps engineer can configure CloudFront to route traffic to the secondary Region when the primary Region is unavailable1
? Create a new origin group. Set the original ALB as the primary origin. Configure
the origin group to fail over for HTTP 5xx status codes. An origin group is a logical grouping of two origins: a primary origin and a secondary origin. By creating an origin group, the DevOps engineer can specify which origin CloudFront should use as a fallback when the primary origin fails. The DevOps engineer can also define which HTTP status codes should trigger a failover from the primary origin to the secondary origin. By setting the original ALB as the primary origin and configuring the origin group to fail over for HTTP 5xx status codes, the DevOps engineer can ensure that CloudFront will switch to the secondary ALB when the primary ALB returns server errors2
? Update the default behavior to use the origin group. A behavior is a set of rules
that CloudFront applies when it receives requests for specific URLs or file types. The default behavior applies to all requests that do not match any other behaviors. By updating the default behavior to use the origin group, the DevOps engineer can enable failover routing for all requests that are sent to the distribution3
This solution will meet the requirements because it will automate the failover of the
application to the secondary Region with zero-second RTO. When CloudFront receives an HTTP GET request, it will first try to route it to the primary ALB in the primary Region. If the primary ALB is healthy and returns a successful response, CloudFront will deliver it to the viewer. If the primary ALB is unhealthy or returns an HTTP 5xx status code, CloudFront will automatically route the request to the secondary ALB in the secondary Region and deliver its response to the viewer.
The other options are not correct because they either do not provide zero-second RTO or do not work as expected. Creating a second CloudFront distribution that has the secondary ALB as the default origin and creating Amazon Route 53 alias records that have a failover policy is not a good option because it will introduce additional latency and complexity to the solution. Route 53 health checks and DNS propagation can take several minutes or longer, which means that viewers might experience delays or errors when accessing the application during a failover event. Creating Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALBs and setting the TTL of both records to O is not a valid option because it will not work with CloudFront distributions. Route 53 does not support health checks for alias records that point to CloudFront distributions, so it cannot detect if an ALB behind a distribution is healthy or not. Creating a CloudFront function that detects HTTP 5xx status codes and returns a 307 Temporary Redirect error response to the secondary ALB is not a valid option because it will not provide zero-second RTO. A 307 Temporary Redirect error response tells viewers to retry their requests with a different URL, which means that viewers will have to make an additional request and wait for another response from CloudFront before reaching the secondary ALB.
References:
? 1: Adding, Editing, and Deleting Origins - Amazon CloudFront
? 2: Configuring Origin Failover - Amazon CloudFront
? 3: Creating or Updating a Cache Behavior - Amazon CloudFront

NEW QUESTION 57
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 nam
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 accoun
D. Associate the template with all the repositorie
E. Add the developer role ARN as an approver.
F. Create an approval rule template for each accoun
G. Associate the template with all repositorie
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 accoun
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 60
A company has 20 service learns 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 65
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 70
A company is running an application on Amazon EC2 instances in an Auto Scaling group. Recently an issue occurred that prevented EC2 instances from launching successfully and it took several hours for the support team to discover the issue. The support team wants to be notified by email whenever an EC2 instance does not start successfully.
Which action will accomplish this?

A. Add a health check to the Auto Scaling group to invoke an AWS Lambda function whenever an instance status is impaired.
B. Configure the Auto Scaling group to send a notification to an Amazon SNS topic whenever a failed instance launch occurs.
C. Create an Amazon CloudWatch alarm that invokes an AWS Lambda function when a failed Attachinstances Auto Scaling API call is made.
D. Create a status check alarm on Amazon EC2 to send a notification to an Amazon SNS topic whenever a status check fail occurs.

Answer: B

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ASGettingNotifications.html#auto-scaling-sns-notifications

NEW QUESTION 75
A company's security policies require the use of security hardened AMIS in production environments. A DevOps engineer has used EC2 Image Builder to create a pipeline that builds the AMIs on a recurring schedule.
The DevOps engineer needs to update the launch templates of the companys Auto Scaling groups. The Auto Scaling groups must use the newest AMIS during the launch of Amazon EC2 instances.
Which solution will meet these requirements with the MOST operational efficiency?

A. Configure an Amazon EventBridge rule to receive new AMI events from Image Builde
B. Target an AWS Systems Manager Run Command document that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
C. Configure an Amazon EventBridge rule to receive new AMI events from Image Builde
D. Target an AWS Lambda function that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
E. Configure the launch template to use a value from AWS Systems Manager Parameter Store for the AMI I
F. Configure the Image Builder pipeline to update the Parameter Store value with the newest AMI ID.
G. Configure the Image Builder distribution settings to update the launch templates with the newest AMI I
H. Configure the Auto Scaling groups to use the newest version of the launch template.

Answer: C

Explanation:
? The most operationally efficient solution is to use AWS Systems Manager Parameter Store1 to store the AMI ID and reference it in the launch template2. This

way, the launch template does not need to be updated every time a new AMI is created by Image Builder. Instead, the Image Builder pipeline can update the Parameter Store value with the newest AMI ID3, and the Auto Scaling group can launch instances using the latest value from Parameter Store.
? The other solutions require updating the launch template or creating a new version of it every time a new AMI is created, which adds complexity and overhead. Additionally, using EventBridge rules and Lambda functions or Run Command documents introduces additional dependencies and potential points of failure.
References: 1: AWS Systems Manager Parameter Store 2: Using AWS Systems Manager parameters instead of AMI IDs in launch templates 3: Update an SSM parameter with
Image Builder

NEW QUESTION 79
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 83
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 Confi
B. Enable trusted access for AWS Config in the organization.
C. Configure a delegated administrator account for AWS Confi
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 aggregato
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 accoun
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 accoun
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 87
A company wants to deploy a workload on several hundred Amazon EC2 instances. The company will provision the EC2 instances in an Auto Scaling group by using a launch template.
The workload will pull files from an Amazon S3 bucket, process the data, and put the results into a different S3 bucket. The EC2 instances must have least-privilege permissions and must use temporary security credentials.
Which combination of steps will meet these requirements? (Select TWO.)

A. Create an IAM role that has the appropriate permissions for S3 bucket
B. Add the IAM role to an instance profile.
C. Update the launch template to include the IAM instance profile.
D. Create an IAM user that has the appropriate permissions for Amazon S3. Generate a secret key and token.
E. Create a trust anchor and profil
F. Attach the IAM role to the profile.
G. Update the launch templat
H. Modify the user data to use the new secret key and token.

Answer: AB

Explanation:
To meet the requirements of deploying a workload on several hundred EC2 instances with least-privilege permissions and temporary security credentials, the company should use an IAM role and an instance profile. An IAM role is a way to grant permissions to an entity that you trust, such as an EC2 instance. An instance profile is a container for an IAM role that you can use to pass role information to an EC2 instance when the instance starts. By using an IAM role and an instance profile, the EC2 instances can automatically receive temporary security credentials from the AWS Security Token Service (STS) and use them to access the S3 buckets. This way, the company does not need to manage or rotate any long-term credentials, such as IAM users or access keys.
To use an IAM role and an instance profile, the company should create an IAM role that has the appropriate permissions for S3 buckets. The permissions should allow the EC2 instances to read from the source S3 bucket and write to the destination S3 bucket. The company should also create a trust policy for the IAM role that specifies that EC2 is allowed to assume the role. Then, the company should add the IAM role to an instance profile. An instance profile can have only one IAM role, so the company does not need to create
multiple roles or profiles for this scenario.
Next, the company should update the launch template to include the IAM instance profile. A launch template is a way to save launch parameters for EC2 instances, such as the instance type, security group, user data, and IAM instance profile. By using a launch template, the company can ensure that all EC2 instances in the Auto Scaling group have consistent configuration and permissions. The company should specify the name or ARN of the IAM instance profile in the launch template. This way, when the Auto Scaling group launches new EC2 instances based on the launch template, they will automatically receive the IAM role and its permissions through the instance profile.
The other options are not correct because they do not meet the requirements or follow best practices. Creating an IAM user and generating a secret key and token is not a good option because it involves managing long-term credentials that need to be rotated regularly. Moreover, embedding credentials in user data is not secure because user data is visible to anyone who can describe the EC2 instance. Creating a trust anchor and profile is not a valid option because trust anchors are used for certificate-based authentication, not for IAM roles or instance profiles. Modifying user data to use a new secret key and token is also not a good option because it requires updating user data every time the credentials change, which is not scalable or efficient.
References:
? 1: AWS Certified DevOps Engineer - Professional Certification | AWS Certification
| AWS
? 2: DevOps Resources - Amazon Web Services (AWS)
? 3: Exam Readiness: AWS Certified DevOps Engineer - Professional
? : IAM Roles for Amazon EC2 - AWS Identity and Access Management
? : Working with Instance Profiles - AWS Identity and Access Management
? : Launching an Instance Using a Launch Template - Amazon Elastic Compute Cloud
? : Temporary Security Credentials - AWS Identity and Access Management

NEW QUESTION 89
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 organizatio
D. Create an SCP to prevent users from modifying VPC flow logs.
E. Turn on AWS Confi
F. Create an AWS Config rule to check whether VPC flow logs are turned o
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 log
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 92
A DevOps engineer is building a continuous deployment pipeline for a serverless application that uses AWS Lambda functions. The company wants to reduce the customer impact of an unsuccessful deployment. The company also wants to monitor for issues.
Which deploy stage configuration will meet these requirements?

A. Use an AWS Serverless Application Model (AWS SAM) template to define the serverless applicatio
B. Use AWS CodeDeploy to deploy the Lambda functions with the Canary10Percent15Minutes Deployment Preference Typ
C. Use Amazon CloudWatch alarms to monitor the health of the functions.
D. Use AWS CloudFormation to publish a new stack update, and include Amazon CloudWatch alarms on all resource
E. Set up an AWS CodePipeline approval action for a developer to verify and approve the AWS CloudFormation change set.
F. Use AWS CloudFormation to publish a new version on every stack update, and include Amazon CloudWatch alarms on all resource
G. Use the RoutingConfig property of the AWS::Lambda::Alias resource to update the traffic routing during the stack update.
H. Use AWS CodeBuild to add sample event payloads for testing to the Lambda function
I. Publish a new version of the functions, and include Amazon CloudWatch alarm
J. Update the production alias to point to the new versio
K. Configure rollbacks to occur when an alarm is in the ALARM state.

Answer: D

Explanation:
 Use routing configuration on an alias to send a portion of traffic to a second function version. For example, you can reduce the risk of deploying a new version by configuring the alias to send most of the traffic to the existing version, and only a small percentage of traffic to the new version.
https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html
The following are the steps involved in the deploy stage configuration that will meet the requirements:
? Use AWS CodeBuild to add sample event payloads for testing to the Lambda
functions.
? Publish a new version of the functions, and include Amazon CloudWatch alarms.
? Update the production alias to point to the new version.
? Configure rollbacks to occur when an alarm is in the ALARM state.
This configuration will help to reduce the customer impact of an unsuccessful deployment
by deploying the new version of the functions to a staging environment first. This will allow the DevOps engineer to test the new version of the functions before deploying it to production.
The configuration will also help to monitor for issues by including Amazon CloudWatch alarms. These alarms will alert the DevOps engineer if there are any problems with the new version of the functions.

NEW QUESTION 94
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 togs 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 96
A rapidly growing company wants to scale for developer demand for AWS development environments. Development environments are created manually in the AWS Management Console. The networking team uses AWS CloudFormation to manage the networking infrastructure, exporting stack output values for the Amazon VPC and all subnets. The development environments have common standards, such as Application Load Balancers, Amazon EC2 Auto Scaling groups, security groups, and Amazon DynamoDB tables.
To keep up with demand, the DevOps engineer wants to automate the creation of development environments. Because the infrastructure required to support the application is expected to grow, there must be a way to easily update the deployed infrastructure. CloudFormation will be used to create a template for the development environments.
Which approach will meet these requirements and quickly provide consistent AWS environments for developers?

A. Use Fn::ImportValue intrinsic functions in the Resources section of the template to retrieve Virtual Private Cloud (VPC) and subnet value

B. Use CloudFormation StackSets for the development environments, using the Count input parameter to indicate the number of environments neede
C. Use the UpdateStackSet command to update existing development environments.
D. Use nested stacks to define common infrastructure component
E. To access the exported values, use TemplateURL to reference the networking team's templat
F. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root templat
G. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
H. Use nested stacks to define common infrastructure component
I. Use Fn::ImportValue intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet value
J. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
K. Use Fn::ImportValue intrinsic functions in the Parameters section of the root template to retrieve Virtual Private Cloud (VPC) and subnet value
L. Define the development resources in the order they need to be created in the CloudFormation nested stack
M. Use the CreateChangeSe
N. and ExecuteChangeSet commands to update existing development environments.

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function- reference-importvalue.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-importvalue.html CF of network exports the VPC, subnet or needed information CF of application imports the above information to its stack and UpdateChangeSet/ ExecuteChangeSet

NEW QUESTION 99
A DevOps engineer has implemented a CI/CO pipeline to deploy an AWS Cloud Format ion template that provisions a web application. The web application consists of an Application Load Balancer (ALB) a target group, a launch template that uses an Amazon Linux 2 AMI an Auto Scaling group of Amazon EC2 instances, a security group and an Amazon RDS for MySQL database The launch template includes user data that specifies a script to install and start the application.
The initial deployment of the application was successful. The DevOps engineer made changes to update the version of the application with the user data. The CI/CD pipeline has deployed a new version of the template However, the health checks on the ALB are now failing The health checks have marked all targets as unhealthy.
During investigation the DevOps engineer notices that the Cloud Formation stack has a status of UPDATE_COMPLETE. However, when the DevOps engineer connects to one of the EC2 instances and checks /varar/log messages, the DevOps engineer notices that the Apache web server failed to start successfully because of a configuration error
How can the DevOps engineer ensure that the CloudFormation deployment will fail if the user data fails to successfully finish running?

A. Use the cfn-signal helper script to signal success or failure to CloudFormation Use the WaitOnResourceSignals update policy within the CloudFormation template Set an appropriate timeout for the update policy.
B. Create an Amazon CloudWatch alarm for the UnhealthyHostCount metri
C. Include an appropriate alarm threshold for the target group Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation
D. Create a lifecycle hook on the Auto Scaling group by using the AWS AutoScaling LifecycleHook resource Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation Set an appropriate timeout on the lifecycle hook.
E. Use the Amazon CloudWatch agent to stream the cloud-init logs Create a subscription filter that includes an AWS Lambda function with an appropriate invocation timeout Configure the Lambda function to use the SignalResource API operation to signal success or failure to CloudFormation.

Answer: A

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-updatepolicy.html

NEW QUESTION 103
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 106
An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and

maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files m source control.
Which solution will accomplish this?

A. In the CloudFormaiion template add an AWS Config rul
B. Place the configuration file content in the rule's InputParameters property and set the Scope property to the EC2 Auto Scaling grou
C. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
D. In the CloudFormation template add an EC2 launch template resourc
E. Place the configuration file content in the launch templat
F. Configure the cfn-mit script to run when the instance is launched and configure the cfn-hup script to poll for updates to the configuration.
G. In the CloudFormation template add an EC2 launch template resourc
H. Place the configuration file content in the launch templat
I. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
J. In the CloudFormation template add CloudFormation imt metadat
K. Place the configuration file content m the metadat
L. Configure the cfn-init script to run when the instance is launched and configure the cfn-hup script to poll for updates to the configuration.

Answer: D

Explanation:
Use the AWS::CloudFormation::Init type to include metadata on an Amazon EC2 instance for the cfn-init helper script. If your template calls the cfn-init script, the script looks for resource metadata rooted in the AWS::CloudFormation::Init metadata key. Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws- resource-init.html

NEW QUESTION 109
A company updated the AWS Cloud Formation 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 AWSC loud Formation FullAccess IAM policy to the AWS CtoudFormation 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 112
The security team depends on AWS CloudTrail to detect sensitive security issues in the company's AWS account. The DevOps engineer needs a solution to auto-remediate CloudTrail being turned off in an AWS account.
What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries?

A. Create an Amazon EventBridge rule for the CloudTrail StopLogging even
B. Create an AWS Lambda (unction that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was calle
C. Add the Lambda function ARN as a target to the EventBridge rule.
D. Deploy the AWS-managed CloudTrail-enabled AWS Config rule set with a periodic interval to 1 hou
E. Create an Amazon EventBridge rule tor AWS Config rules compliance chang
F. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLoggmg was calle
G. Add the Lambda function ARN as a target to the EventBridge rule.
H. Create an Amazon EventBridge rule for a scheduled event every 5 minute
I. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on a CloudTrail trail in the AWS accoun
J. Add the Lambda function ARN as a target to the EventBridge rule.
K. Launch a t2 nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current accoun
L. If the CloudTrail trail is disabled have the script re-enable the trail.

Answer: A

Explanation:
https://aws.amazon.com/blogs/mt/monitor-changes-and-auto-enable-logging-in-aws-cloudtrail/

NEW QUESTION 115
A company has an application that includes AWS Lambda functions. The Lambda functions run Python code that is stored in an AWS CodeCommit repository. The company has recently experienced failures in the production environment because of an error in the Python code. An engineer has written unit tests for the Lambda functions to help avoid releasing any future defects into the production environment.
The company's DevOps team needs to implement a solution to integrate the unit tests into an existing AWS CodePipeline pipeline. The solution must produce reports about the unit tests for the company to view.
Which solution will meet these requirements?

A. Associate the CodeCommit repository with Amazon CodeGuru Reviewe
B. Create a new AWS CodeBuild projec
C. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild projec
D. Create a buildspec.yml file in the CodeCommit repositor

E. In the buildspec.yml file, define the actions to run a CodeGuru review.
F. Create a new AWS CodeBuild projec
G. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild projec
H. Create a CodeBuild report grou
I. Create a buildspec.yml file in the CodeCommit repositor
J. In the buildspec.yml file, define the actions to run the unit tests with an output of JUNITXML in the build phase section.Configure the test reports to be uploaded to the new CodeBuild report group.
K. Create a new AWS CodeArtifact repositor
L. Create a new AWS CodeBuild projec
M. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild projec
N. Create an appspec.yml file in the original CodeCommit repositor
O. In the appspec.yml file, define the actions to run the unit tests with an output of CUCUMBERJSON in the build phase sectio
P. Configure the tests reports to be sent to the new CodeArtifact repository.
Q. Create a new AWS CodeBuild projec
R. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild projec
S. Create a new Amazon S3 bucke
T. Create a buildspec.yml file in the CodeCommit repositor
. In the buildspec.yml file, define the actions to run the unit tests with an output of HTML in the phases sectio
. In the reports section, upload the test reports to the S3 bucket.

Answer: B

Explanation:
The correct answer is B. Creating a new AWS CodeBuild project and configuring a test stage in the AWS CodePipeline pipeline that uses the new CodeBuild project is the best way to integrate the unit tests into the existing pipeline. Creating a CodeBuild report group and uploading the test reports to the new CodeBuild report group will produce reports about the unit tests for the company to view. Using JUNITXML as the output format for the unit tests is supported by CodeBuild and will generate a valid report. Option A is incorrect because Amazon CodeGuru Reviewer is a service that provides automated code reviews and recommendations for improving code quality and performance. It is not a tool for running unit tests or producing test reports. Therefore, option A will not meet the requirements.
Option C is incorrect because AWS CodeArtifact is a service that provides secure, scalable, and cost-effective artifact management for software development. It is not a tool for running unit tests or producing test reports. Moreover, option C uses CUCUMBERJSON as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.
Option D is incorrect because uploading the test reports to an Amazon S3 bucket is not the best way to produce reports about the unit tests for the company to view. CodeBuild has a built-in feature to create and manage test reports, which is more convenient and efficient than using S3. Furthermore, option D uses HTML as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.

NEW QUESTION 118
A company is divided into teams Each team has an AWS account and all the accounts are in an organization in AWS Organizations. Each team must retain full administrative rights to its AWS account. Each team also must be allowed to access only AWS services that the company approves for use AWS services must gam approval through a request and approval process.
How should a DevOps engineer configure the accounts to meet these requirements?

A. Use AWS CloudFormation StackSets to provision IAM policies in each account to deny access to restricted AWS service
B. In each account configure AWS Config rules that ensure that the policies are attached to IAM principals in the account.
C. Use AWS Control Tower to provision the accounts into OUs within the organization Configure AWS Control Tower to enable AWS IAM identity Center (AWS Single Sign-On). Configure 1AM Identity Center to provide administrative access Include deny policies on user roles for restricted AWS services.
D. Place all the accounts under a new top-level OU within the organization Create an SCP that denies access to restricted AWS services Attach the SCP to the OU.
E. Create an SCP that allows access to only approved AWS service
F. Attach the SCP to the root OU of the organizatio
G. Remove the FullAWSAccess SCP from the root OU of the organization.

Answer: C

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html A managed prefix list is a set of one or more CIDR blocks. You can use prefix lists to make it easier to configure and maintain your security groups and route tables. https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html With AWS Resource Access Manager (AWS RAM), the owner of a prefix list can share a prefix list with the following: Specific AWS accounts inside or outside of its organization in AWS Organizations An organizational unit inside its organization in AWS Organizations An entire organization in AWS Organizations

NEW QUESTION 121
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 grou
B. Create a warm pool instance in the stopped stat
C. Define the warm pool siz
D. Create a new version of the launch template that has detailed monitoring enable
E. use Spot Instances.
F. For the critical data, modify the existing Auto Scaling grou
G. Create a warm pool instance in the stopped stat
H. Define the warm pool siz
I. Create a new version of the launch template that has detailed monitoring enable
J. Use On-Demand Instances.
K. For the critical dat
L. modify the existing Auto Scaling grou
M. Create a lifecycle hook to ensure that bootstrap scripts are completed successfull

N. Ensure that the application on the instances is ready to accept traffic before the instances are registere
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 templat
Q. Configure the launch template to install the unified Amazon CloudWatch agent and to configure the CloudWatch agent with a custom memory utilization metri
R. Use Spot Instance
S. Add the new Auto Scaling group as the target group for the AL
T. Modify the application to use two target groups for critical data and noncritical data.
. For the noncritical data, create a second Auto Scaling grou
. Choose the predefined memory utilization metric type for the target tracking scaling polic
. Use Spot Instance
. Add the new Auto Scaling group as the target group for the AL
. 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 126
......