AWS-Certified-DevOps-Engineer-Professional Dumps

Amazon AWS Certified DevOps Engineer Professional

https://www.certleader.com/AWS-Certified-DevOps-Engineer-Professional-dumps.html

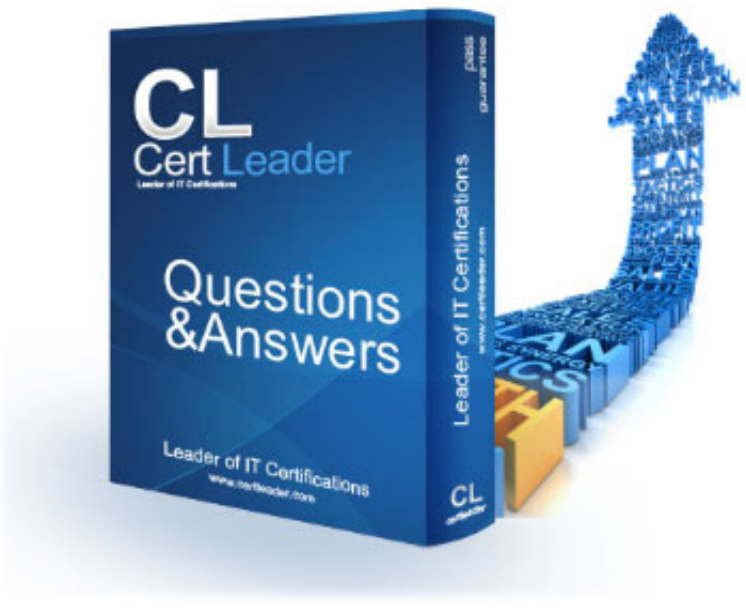

NEW QUESTION 1
A company is using AWS Organizations to create separate AWS accounts for each of its departments. It needs to automate the following tasks:
Updating the Linux AMIs with new patches periodically and generating a golden image Installing a new version of Chef agents in the golden image, if available
Enforcing the use of the newly generated golden AMIs in the department's account Which option requires the LEAST management overhead?

A. Write a script to launch an Amazon EC2 instance from the previous golden AMI, apply the patch updates, install the new version of the Chef agent, generate a new golden AMI, and then modify the AMI permissions to share only the new image with the departments' accounts.
B. Use an AWS Systems Manager Run Command to update the Chef agent first, use Amazon EC2 Systems Manager Automation to generate an updated AMI, and then assume an IAM role to copy the new golden AMI into the departments' accounts.
C. Use AWS Systems Manager Automation to update the Linux AMI using the previous image, provide the URL for the script that will update the Chef agent, and then use AWS Organizations to replace the previous golden AMI into the departments' accounts.
D. Use AWS Systems Manager Automation to update the Linux AMI from the previous golden image, provide the URL for the script that will update the Chef agent, and then share only the newly generated AMI with the departments' accounts.

Answer: C

NEW QUESTION 2
A Security team requires all Amazon EBS volumes that are attached to an Amazon EC2 instance to have AWS Key Management Service (AWS KMS) encryption enabled. If encryption is not enabled, the company's policy requires the EBS volume to be detached and deleted. A DevOps Engineer must automate the detection and deletion of unencrypted EBS volumes. Which method should the Engineer use to accomplish this with the LEAST operational effort?

A. Create an Amazon CloudWatch Events rule that invokes an AWS Lambda function when an EBS volume is create
B. The Lambda function checks the EBS volume for encryptio
C. If encryption is not enabled and the volume is attached to an instance, the function deletes the volume.
D. Create an AWS Lambda function to describe all EBS volumes in the region and identify volumes that are attached to an EC2 instance without encryption enable
E. The function then deletes all non-compliant volume
F. The AWS Lambda function is invoked every 5 minutes by an Amazon CloudWatch Events scheduled rule.
G. Create a rule in AWS Config to check for unencrypted and attached EBS volume
H. Subscribe an AWS Lambda function to the Amazon SNS topic that AWS Config sends change notifications t
I. The Lambda function checks the change notification and deletes any EBS volumes that are non-compliant.
J. Launch an EC2 instance with an IAM role that has permissions to describe and delete volume
K. Run ascript on the EC2 instance every 5 minutes to describe all EBS volumes in all regions and identify volumes that are attached without encryption enable
L. The script then deletes those volumes.

Answer: B

NEW QUESTION 3
A highly regulated company has a policy that DevOps Engineers should not log in to their Amazon EC2 instances except in emergencies. If a DevOps Engineer does log in, the Security team must be notified within 15 minutes of the occurrence.
Which solution will meet these requirements?

A. Install the Amazon Inspector agent on each EC2 instanc
B. Subscribe to Amazon CloudWatch Events notification
C. Trigger an AWS Lambda function to check if a message is about user login
D. If it is, send a notification to the Security team using Amazon SNS.
E. Install the Amazon CloudWatch agent on each EC2 instanc
F. Configure the agent to push all logs to Amazon CloudWatch Logs and set up a CloudWatch metric filter that searches for user login
G. If a login is found, send a notification to the Security team using Amazon SNS.
H. Set up AWS CloudTrail with Amazon CloudWatch Log
I. Subscribe CloudWatch Logs to Amazon Kinesi
J. Attach AWS Lambda to Kinesis to parse and determine if a log contains a user logi
K. If it does, send a notification to the Security team using Amazon SNS.
L. Set up a script on each Amazon EC2 instance to push all logs to Amazon S3. Set up an S3 event to trigger an AWS Lambda function, which triggers an Amazon Athena query to ru
M. The Athena query checks for logins and sends the output to the Security team using Amazon SNS.

Answer: B

NEW QUESTION 4
You have deployed an application to AWS which makes use of Autoscaling to launch new instances. You now want to change the instance type for the new instances. Which of the following is one of the action items to achieve this deployment?

A. Use Elastic Beanstalk to deploy the new application with the new instance type
B. Use Cloudformation to deploy the new application with the new instance type
C. Create a new launch configuration with the new instance type
D. Create new EC2 instances with the new instance type and attach it to the Autoscaling Group

Answer: C

Explanation:
The ideal way is to create a new launch configuration, attach it to the existing Auto Scaling group, and terminate the running instances.
Option A is invalid because Clastic beanstalk cannot launch new instances on demand. Since the current scenario requires Autoscaling, this is not the ideal option
Option B is invalid because this will be a maintenance overhead, since you just have an Autoscaling Group.
There is no need to create a whole Cloudformation
template for this.
Option D is invalid because Autoscaling Group will still launch CC2 instances with the older launch configuration
For more information on Autoscaling Launch configuration, please refer to the below document link: from AWS

http://docs.aws.amazon.com/autoscaling/latest/userguide/l_aunchConfiguration.html

NEW QUESTION 5
A company is building a web and mobile application that uses a serverless architecture powered by AWS Lambda and Amazon API Gateway. The company wants to fully automate the backend Lambda deployment based on code that is pushed to the appropriate environment branch in an AWS CodeCommit repository.
The deployment must have the following:
*Separate environment pipelines for testing and production.
*Automatic deployment that occurs for test environments only. Which steps should be taken to meet these requirements?

A. Configure a new AWS CodePipeline servic
B. Create a CodeCommit repository for each environment.Set up CodePipeline to retrieve the source code from the appropriate repositor
C. Set up a deployment step to deploy the Lambda functions with AWS CloudFormation.
D. Create two AWS CodePipeline configurations for test and production environment
E. Configure the production pipeline to have a manual approval ste
F. Create a CodeCommit repository for each environmen
G. Set up each CodePipeline to retrieve the source code from the appropriate repositor
H. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
I. Create two AWS CodePipeline configurations for test and production environment
J. Configure the production pipeline to have a manual approval ste
K. Create one CodeCommit repository with a branch for each environmen
L. Set up each CodePipeline to retrieve the source code from the appropriate branch in the repositor
M. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
N. Create an AWS CodeBuild configuration for test and production environment
O. Configure the production pipeline to have a manual approval ste
P. Create one CodeCommit repository with a branch for each environmen
Q. Push the Lambda function code to an Amazon S3 bucke
R. Set up the deployment step to deploy the Lambda functions from the S3 bucket.

Answer: B

NEW QUESTION 6
A global company with distributed Development teams built a web application using a microservices architecture running on Amazon ECS. Each application service is independent and runs as a service in the ECS cluster. The container build files and source code reside in a private GitHub source code repository. Separate ECS clusters exist for development, testing, and production environments.
Developers are required to push features to branches in the GitHub repository and then merge the changes into an environment-specific branch (development, test, or production). This merge needs to trigger an automated pipeline to run a build and a deployment to the appropriate ECS cluster.
What should the DevOps Engineer recommend as an automated solution to these requirements?

A. Create an AWS CloudFormation stack for the ECS cluster and AWS CodePipeline service
B. Store the container build files in an Amazon S3 bucke
C. Use a post-commit hook to trigger a CloudFormation stack update that deploys the ECS cluste
D. Add a task in the ECS cluster to build and push images to Amazon ECR, based on the container build files in S3.
E. Create a separate pipeline in AWS CodePipeline for each environmen
F. Trigger each pipeline based on commits to the corresponding environment branch in GitHu
G. Add a build stage to launch AWS CodeBuild to create the container image from the build file and push it to Amazon EC
H. Then add another stage to update the Amazon ECS task and service definitions in the appropriate cluster for that environment.
I. Create a pipeline in AWS CodePipelin
J. Configure it to be triggered by commits to the master branch in GitHu
K. Add a stage to use the Git commit message to determine which environment the commit should be applied to, then call the create-image Amazon ECR command to build the image, passing it to the container build fil
L. Then add a stage to update the ECS task and service definitions in the appropriate cluster for that environment.
M. Create a new repository in AWS CodeCommi
N. Configure a scheduled project in AWS CodeBuild to synchronize the GitHub repository to the new CodeCommit repositor
O. Create a separate pipeline for each environment triggered by changes to the CodeCommit repositor
P. Add a stage using AWS Lambda to build the container image and push to Amazon EC
Q. Then add another stage to update the ECS task and service definitions in the appropriate cluster for that environment.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-cd-pipeline.html

NEW QUESTION 7
A consulting company was hired to assess security vulnerabilities within a client company's application and propose a plan to remediate all identified issues. The architecture is identified as follows: Amazon S3 storage for content, an Auto Scaling group of Amazon EC2 instances behind an Elastic Load Balancer with attached Amazon EBS storage, and an Amazon RDS MySQL database. There are also several AWS Lambda functions that communicate directly with the RDS database using connection string statements in the code.
The consultants identified the top security threat as follows: the application is not meeting its requirement to have encryption at rest.
What solution will address this issue with the LEAST operational overhead and will provide monitoring for
potential future violations?

A. Enable SSE encryption on the S3 buckets and RDS databas
B. Enable OS-based encryption of data on EBS volume
C. Configure Amazon Inspector agents on EC2 instances to report on insecure encryption cipher
D. Set up AWS Config rules to periodically check for non-encrypted S3 objects.
E. Configure the application to encrypt each file prior to storing on Amazon S3. Enable OS-based encryption of data on EBS volume
F. Encrypt data on write to RD
G. Run cron jobs on each instance to check for encrypted data and notify via Amazon SN
H. Use S3 Events to call an AWS Lambda function and verify if the file is encrypted.

I. Enable Secure Sockets Layer (SSL) on the load balancer, ensure that AWS Lambda is using SSL to communicate to the RDS database, and enable S3 encryptio
J. Configure the application to force SSL for incoming connections and configure RDS to only grant access if the session is encrypte
K. Configure Amazon Inspector agents on EC2 instances to report on insecure encryption ciphers.
L. Enable SSE encryption on the S3 buckets, EBS volumes, and the RDS databas
M. Store RDS credentials in EC2 Parameter Stor
N. Enable a policy on the S3 bucket to deny unencrypted put
O. Set up AWS Config rules to periodically check for non-encrypted S3 objects and EBS volumes, and to ensure that RDS storage is encrypted.

Answer: D

NEW QUESTION 8
A company recently launched an application that is more popular than expected. The company wants to ensure the application can scale to meet increasing demands and provide reliability using multiple Availability Zones (AZs) The application runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) A DevOps engineer has created an Auto Scaling group across multiple AZs for the application Instances launched in the newly added AZs are not receiving any traffic for the application.
What is likely causing this issue?

A. Auto Scaling groups can create new instances in a single AZ only.
B. The EC2 instances have not been manually associated to the ALB
C. The ALB should be replaced with a Network Load Balancer (NLB).
D. The new AZ has not been added to the ALB

Answer: A

NEW QUESTION 9
A government agency is storing highly confidential files in an encrypted Amazon S3 bucket. The agency has configured federated access and has allowed only a particular on-premises Active Directory user group to access this bucket.
The agency wants to maintain audit records and automatically detect and revert any accidental changes administrators make to the IAM policies used for providing this restricted federated access.
Which of the following options provide the FASTEST way to meet these requirements?

A. Configure an Amazon CloudWatch Events Event Bus on an AWS CloudTrail API for triggering the AWS Lambda function that detects and reverts the change.
B. Configure an AWS Config rule to detect the configuration change and execute an AWS Lambda function to revert the change.
C. Schedule an AWS Lambda function that will scan the IAM policy attached to the federated access role for detecting and reverting any changes.
D. Restrict administrators in the on-premises Active Directory from changing the IAM policies

Answer: B

Explanation:
https://www.puresec.io/blog/aws-security-best-practices-config-rules-lambda-security "Cloudwatch Event Bus" are used for -> "Sending and Receiving Events Between AWS Accounts"
https://aws.amazon.com/about-aws/whats-new/2017/06/cloudwatch-events-adds-cross-account-event-delivery-s
https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config-rules.html

NEW QUESTION 10
An Application team has three environments for their application: development, pre-production, and production. The team recently adopted AWS CodePipeline. However, the team has had several deployments of misconfigured or nonfunctional development code into the production environment, resulting in user disruption and downtime. The DevOps Engineer must review the pipeline and add steps to identify problems with the application before it is deployed.
What should the Engineer do to identify functional issues during the deployment process? (Choose two.)

A. Use Amazon Inspector to add a test action to the pipelin
B. Use the Amazon Inspector Runtime Behavior Analysis Inspector rules package to check that the deployed code complies with company security standards before deploying it to production.
C. Using AWS CodeBuild to add a test action to the pipeline to replicate common user activities and ensure that the results are as expected before progressing to production deployment.
D. Create an AWS CodeDeploy action in the pipeline with a deployment configuration that automatically deploys the application code to a limited number of instance
E. The action then pauses the deployment so that the QA team can review the application functionalit
F. When the review is complete, CodeDeploy resumes and deploys the application to the remaining production Amazon EC2 instances.
G. After the deployment process is complete, run a testing activity on an Amazon EC2 instance in a different region that accesses the application to simulate user behavio
H. If unexpected results occur, the testing activity sends a warning to an Amazon SNS topi
I. Subscribe to the topic to get updates.
J. Add an AWS CodeDeploy action in the pipeline to deploy the latest version of the development code to pre-productio
K. Add a manual approval action in the pipeline so that the QA team can test and confirm the expected functionalit
L. After the manual approval action, add a second CodeDeploy action that deploys the approved code to the production environment.

Answer: BE

Explanation:
https://docs.aws.amazon.com/codepipeline/latest/userguide/integrations-action-type.html#integrations-test
https://docs.aws.amazon.com/codepipeline/latest/userguide/integrations-action-type.html#integrations-deploy

NEW QUESTION 10
A company wants to migrate a legacy application to AWS and develop a deployment pipeline that uses AWS services only. A DevOps engineer is migrating all of the application code from a Git repository to AWS CodeCommit while preserving the history of the repository. The DevOps engineer has set all the permissions within CodeCommit, installed the Git client and the AWS CLI on a local computer, and is ready to migrate the repository.
Which actions will follow?

A. Create the CodeCommit repository using the AWS CL
B. Clone the Git repository directly to CodeCommit using the AWS CL
C. Validate that the files were migrated, and publish the CodeCommit repository.
D. Create the CodeCommit repository using the AWS Management Consol
E. Clone both the Git and CodeCommit repositories to the local compute
F. Copy the files from the Git repository to the CodeCommit repository on the local compute
G. Commit the CodeCommit repertor
H. Validate that the files were migrated, and share the CodeCommit repository.
I. Create the CodeCommit repository using the AWS Management Consol
J. Use the console to clone the Git repository into the CodeCommit repositor
K. Validate that the files were migrated, and publish the CodeCommit repository.
L. Create the CodeCommit repository using the AWS Management Console or the AWS CL
M. Clone the Git repository with a mirror argument to the local computer and push the repository to CodeCommi
N. Validate that the files were migrated, and share the CodeCommit repository.

Answer: D

NEW QUESTION 12
A company has an application that has predictable peak traffic times. The company wants the application instances to scale up only during the peak times. The application stores state in Amazon DynamoDB. The application environment uses a standard Node.js application stack and custom Chef recipes stored in a private Git repository.
Which solution is MOST cost-effective and requires the LEAST amount of management overhead when performing rolling updates of the application environment?

A. Create a custom AMI with the Node.js environment and application stack using Chef recipe
B. Use the AMI in an Auto Scaling group and set up scheduled scaling for the required times, then set up an Amazon EC2 IAM role that provides permission to access DynamoDB.
C. Create a Docker file that uses the Chef recipes for the application environment based on an official Node.js Docker imag
D. Create an Amazon ECS cluster and a service for the application environment, then create a task based on this Docker imag
E. Use scheduled scaling to scale the containers at the appropriate times and attach a task-level IAM role that provides permission to access DynamoDB.
F. Configure AWS OpsWorks stacks and use custom Chef cookbook
G. Add the Git repository information where the custom recipes are stored, and add a layer in OpsWorks for the Node.js application serve
H. Then configure the custom recipe to deploy the application in the deploy ste
I. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB.
J. Configure AWS OpsWorks stacks and push the custom recipes to an Amazon S3 bucket and configure custom recipes to point to the S3 bucke
K. Then add an application layer type for a standard Node.js application server and configure the custom recipe to deploy the application in the deploy step from the S3 bucke
L. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB

Answer: D

NEW QUESTION 15
A production account has a requirement that any Amazon EC2 instance that has been logged into manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with Amazon CloudWatch Logs agent configured.
How can this process be automated?

A. Create a CloudWatch Logs subscription to an AWS Step Functions applicatio
B. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissione
C. Then create a CloudWatch Events rule to trigger a second AWS Lambda function once a day that will terminate all instances with this tag.
D. Create a CloudWatch alarm that will trigger on the login even
E. Send the notification to an Amazon SNS topic that the Operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.
F. Create a CloudWatch alarm that will trigger on the login even
G. Configure the alarm to send to an Amazon SQS queu
H. Use a group of worker instances to process messages from the queue, which then schedules the Amazon CloudWatch Events rule to trigger.
I. Create a CloudWatch Logs subscription in an AWS Lambda functio
J. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissione
K. Create a CloudWatch Events rule to trigger a daily Lambda function that terminates all instances withthis ta

Answer: D

Explanation:
https://boto3.amazonaws.com/v1/documentation/api/latest/guide/cw-example-subscription-filters.html

NEW QUESTION 20
An application has microservices spread across different AWS accounts and is integrated with an on-premises legacy system for some of its functionality. Because of the segmented architecture and missing logs, every time the application experiences issues, it is taking too long to gather the logs to identify the issues. A DevOps Engineer must fix the log aggregation process and provide a way to centrally analyze the logs.
Which is the MOST efficient and cost-effective solution?

A. Collect system logs and application logs by using the Amazon CloudWatch Logs agen
B. Use the Amazon S3 API to export on-premises logs, and store the logs in an S3 bucket in a central accoun
C. Build an Amazon EMR cluster to reduce the logs and derive the root cause.
D. Collect system logs and application logs by using the Amazon CloudWatch Logs agen
E. Use the Amazon S3 API to import on-premises log
F. Store all logs in S3 buckets in individual account
G. Use Amazon Macie to write a query to search for the required specific event-related data point.
H. Collect system logs and application logs using the Amazon CloudWatch Logs agen
I. Install the CloudWatch Logs agent on the on-premises server
J. Transfer all logs from AWS to the on-premises data cente
K. Use an Amazon Elasticsearch Logstash Kibana stack to analyze logs on premises.
L. Collect system logs and application logs by using the Amazon CloudWatch Logs agen

M. Install a CloudWatch Logs agent for on-premises resource
N. Store all logs in an S3 bucket in a central accoun
O. Set up an Amazon S3 trigger and an AWS Lambda function to analyze incoming logs and automatically identify anomalie
P. Use Amazon Athena to run ad hoc queries on the logs in the central account.

Answer: D

NEW QUESTION 24
A DevOps Engineer manages a large commercial website that runs on Amazon EC2. The website uses Amazon Kinesis Data Streams to collect and process web logs. The Engineer manages the Kinesis consumer application, which also runs on EC2. Spikes of data cause the Kinesis consumer application to fall behind, and the streams drop records before they can be processed.
What is the FASTEST method to improve stream handling?

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

Answer: B

NEW QUESTION 29
A DevOps Engineer encountered the following error when attempting to use an AWS CloudFormation template to create an Amazon ECS cluster:
An error occurred (InsufficientCapabilitiesException) when calling the CreateStack operation.
What caused this error and what steps need to be taken to allow the Engineer to successfully execute the AWS CloudFormation template?

A. The AWS user or role attempting to execute the CloudFormation template does not have the permissions required to create the resources within the templat
B. The Engineer must review the user policies and add any permissions needed to create the resources and then rerun the template execution.
C. The AWS CloudFormation service cannot be reached and is not capable of creating the cluste
D. The Engineer needs to confirm that routing and firewall rules are not preventing the AWS CloudFormation script from communicating with the AWS service endpoints, and then rerun the template execution.
E. The CloudFormation execution was not granted the capability to create IAM resource
F. The Engineer needs to provide CAPABILITY_IAM and as capabilities in the CloudFormation execution parameters or provide the capabilities in the AWS Management Consol
G. CAPABILITY_NAMED_IAM
H. CloudFormation is not capable of fulfilling the request of the specified resources in the current AWS Regio
I. The Engineer needs to specify a new region and rerun the template

Answer: C

NEW QUESTION 32
A DevOps Engineer is launching a new application that will be deployed using Amazon Route 53, an Application Load Balancer, Auto Scaling, and Amazon DynamoDB. One of the key requirements of this launch is that the application must be able to scale to meet a sudden load increase. During periods of low usage, the infrastructure components must scale down to optimize cost.
What steps can the DevOps Engineer take to meet the requirements? (Select TWO.)

A. Use AWS Trusted Advisor to submit limit increase requests for the Amazon EC2 instances that will be used by the infrastructure.
B. Determine which Amazon EC2 instance limits need to be raised by leveraging AWS Trusted Advisor, and submit a request to AWS Support to increase those limits.
C. Enable Auto Scaling for the DynamoDB tables that are used by the application.
D. Configure the Application Load Balancer to automatically adjust the target group based on the current load.
E. Create an Amazon CloudWatch Events scheduled rule that runs every 5 minutes to track the current use of the Auto Scaling grou
F. If usage has changed, trigger a scale-up event to adjust the capacit
G. Do the same for DynamoDB read and write capacities.

Answer: BC

NEW QUESTION 37
The Development team at an online retailer has moved to Business support and want to take advantage of the AWS Health Dashboard and the AWS Health API to automate remediation actions for issues with the health of AWS resources. The first use case is to respond to AWS detecting an IAM access key that is listed on a public code repository site. The automated response will be to delete the IAM access key and send a notification to the Security team.
How should this be achieved?

A. Create an AWS Lambda function to delete the IAM access ke
B. Send AWS CloudTrail logs to AWS CloudWatch log
C. Create a CloudWatch Logs metric filter for the AWS_RISK_CREDENTIALS_EXPOSED event with two actions: first, run the Lambda function; second, use Amazon SNS to send a notification to the Security team.
D. Create an AWS Lambda function to delete the IAM access ke
E. Create an AWS Config rule for changes to aws.health and the AWS_RISK_CREDENTIALS_EXPOSED event with two actions: first, run the Lambda function; second, use Amazon SNS to send a notification to the Security team.
F. Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security tea
G. Create an AWS Personal Health Dashboard rule for the AWS_RISK_CREDENTIALS_EXPOSED event; set the target of the Personal Health Dashboard rule to Step Functions.
H. Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security tea
I. Create an Amazon CloudWatch Events rule with an aws.health event source and the AWS_RISK_CREDENTIALS_EXPOSED event, set the target of the CloudWatch Events rule to Step Functions.

Answer: A

NEW QUESTION 41
A company has multiple child 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 child accounts using an AWS Lambda function in the master account of the organization.
Which combination of access changes will meet these requirements? (Select THREE.)

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

Answer: ADF

NEW QUESTION 45
To run an application, a DevOps Engineer launches an Amazon EC2 instances with public IP addresses in a public subnet. A user data script obtains the application artifacts and installs them on the instances upon launch. A change to the security classification of the application now requires the instances to run with no access to the Internet. While the instances launch successfully and show as healthy, the application does not seem to be installed.
Which of the following should successfully install the application while complying with the new rule?

A. Launch the instances in a public subnet with Elastic IP addresses attache
B. Once the application is installed and running, run a script to disassociate the Elastic IP addresses afterwards.
C. Set up a NAT gatewa
D. Deploy the EC2 instances to a private subne
E. Update the private subnet's route table to use the NAT gateway as the default route.
F. Publish the application artifacts to an Amazon S3 bucket and create a VPC endpoint for S3. Assign an IAM instance profile to the EC2 instances so they can read the application artifacts from the S3 bucket.
G. Create a security group for the application instances and whitelist only outbound traffic to the artifact repositor
H. Remove the security group rule once the install is complete.

Answer: C

Explanation:
EC2 instances running in private subnets of a VPC can now have controlled access to S3 buckets, objects, and API functions that are in the same region as the VPC. You can use an S3 bucket policy to indicate which VPCs and which VPC Endpoints have access to your S3 buckets 1- https://aws.amazon.com/pt/blogs/aws/new-vpc-endpoint-for-amazon-s3/

NEW QUESTION 48
An application is running on Amazon EC2. It has an attached IAM role that is receiving an AccessDenied error while trying to access a SecureString parameter resource in the AWS Systems Manager Parameter Store. The SecureString parameter is encrypted with a customer-managed Customer Master Key (CMK),
What steps should the DevOps Engineer take to grant access to the role while granting least privilege? (Select three.)

A. Set ssm:GetParamter for the parameter resource in the instance role's IAM policy.
B. Set kms:Decrypt for the instance role in the customer-managed CMK policy.
C. Set kms:Decrypt for the customer-managed CMK resource in the role's IAM policy.
D. Set ssm:DecryptParameter for the parameter resource in the instance role IAM policy.
E. Set kms:GenerateDataKey for the user on the AWS managed SSM KMS key.
F. Set kms:Decrypt for the parameter resource in the customer-managed CMK policy.

Answer: ABC

NEW QUESTION 52
A government agency has multiple AWS accounts, many of which store sensitive citizen information. A Security team wants to detect anomalous account and network activities (such as SSH brute force attacks) in any account and centralize that information in a dedicated security account. Event information should be stored in an Amazon S3 bucket in the security account, which is monitored by the department's Security Information and Even Manager (SIEM) system.
How can this be accomplished?

A. Enable Amazon Macie in every accoun
B. Configure the security account as the Macie Administrator for every member account using invitation/acceptanc
C. Create an Amazon CloudWatch Events rule in the security account to send all findings to Amazon Kinesis Data Firehouse, which should push the findings to the S3 bucket.
D. Enable Amazon Macie in the security account onl
E. Configure the security account as the Macie Administrator for every member account using invitation/ acceptanc
F. Create an Amazon CloudWatch Events rule in the security account to send all findings to Amazon Kinesis Data Stream
G. Write and application using KCL to read data from the Kinesis Data Streams and write to the S3 bucket.
H. Enable Amazon GuardDuty in every accoun
I. Configure the security account as the GuardDuty Administrator for every member account using invitation/ acceptanc
J. Create an Amazon CloudWatch rule in the security account to send all findings to Amazon Kinesis Data Firehouse, which will push the findings to the S3 bucket.
K. Enable Amazon GuardDuty in the security account onl
L. Configure the security account as the GuardDuty Administrator for every member account using invitation/acceptanc
M. Create an Amazon CloudWatch rule in the security account to send all findings to Amazon Kinesis Data Stream
N. Write and application using KCL to read data from Kinesis Data Streams and write to the S3 bucket.

Answer: C

Explanation:

https://aws.amazon.com/blogs/security/how-to-manage-amazon-guardduty-security-findings-across-multiple-acc

NEW QUESTION 53
A company is deploying a new mobile game on AWS for its customers around the world. The Development team uses AWS Code services and must meet the following requirements:
- Clients need to send/receive real-time playing data from the backend frequently and with minimal latency
- Game data must meet the data residency requirement
Which strategy can a DevOps Engineer implement to meet their needs?

A. Deploy the backend application to multiple region
B. Any update to the code repository triggers a two-stage build and deployment pipelin
C. A successful deployment in one region invokes an AWSLambda function to copy the build artifacts to an Amazon S3 bucket in another regio
D. After the artifactis copied, it triggers a deployment pipeline in the new region.
E. Deploy the backend application to multiple Availability Zones in a single regio
F. Create an Amazon CloudFront distribution to serve the application backend to global customer
G. Any update to the code repository triggers a two-stage build-and-deployment pipelin
H. The pipeline deploys the backend application to all Availability Zones.
I. Deploy the backend application to multiple region
J. Use AWS Direct Connect to serve the application backend to global customer
K. Any update to the code repository triggers a two-stagebuild-and-deployment pipeline in the regio
L. After a successful deployment in the region, the pipeline continues to deploy the artifact to another region.
M. Deploy the backend application to multiple region
N. Any update to the code repository triggers atwo-stage build-and-deployment pipeline in the regio
O. After a successful deployment in the region, the pipeline invokes the pipeline in another region and passes the build artifact locatio
P. The pipeline uses the artifact location and deploys applications in the new region.

Answer: A

NEW QUESTION 55
A company is using an AWS CodeBuild project to build and package an application. The packages are copied to a shared Amazon S3 bucket before being deployed across multiple AWS accounts.
The buildspec.yml file contains the following:

```
version: 0.2
phases:
  build:
    commands:
      - go build -o myapp
  post_build:
    commands:
      - aws s3 cp --acl authenticated-read myapp s3://artifacts/
```

The DevOps Engineer has noticed that anybody with an AWS account is able to download the artifacts. What steps should the DevOps Engineer take to stop this?

A. Modify the post_build to command to use ""-acl public-read and configure a bucket policy that grants read access to the relevant AWS accounts only.
B. Configure a default ACL for the S3 bucket that defines the set of authenticated users as the relevant AWS accounts only and grants read-only access.
C. Create an S3 bucket policy that grants read access to the relevant AWS accounts and denies read access to the principal "*"
D. Modify the post_build command to remove ""-acl authenticated-read and configure a bucket policy that allows read access to the relevant AWS accounts only.

Answer: D

NEW QUESTION 58
You are responsible for your company's large multi-tiered Windows-based web application running on Amazon EC2 instances situated behind a load balancer. While reviewing metrics, you've started noticing an upwards trend for slow customer page load time. Your manager has asked you to come up with a solution to ensure that customer load time is not affected by too many requests per second. Which technique would you use to solve this issue?

A. Re-deploy your infrastructure usingan AWS CloudFormation templat
B. Configure Elastic Load Balancing health checks to initiate a new AWS CloudFormation stack when health checks return failed.
C. Re-deploy your infrastructure using an AWS CloudFormation templat
D. Spin up a second AWS CloudFormation stac
E. Configure Elastic Load Balancing SpillOver functionality to spill over any slow connections to the second AWS CloudFormation stack.
F. Re-deploy your infrastructure using AWS CloudFormation, Elastic Beanstalk, and Auto Scalin
G. Set up your Auto Scalinggroup policies to scale based on the number of requests per second as well as the current customer load tim
H. •>/D- Re-deploy your application using an Auto Scaling templat
I. Configure the Auto Scaling template to spin up a new Elastic Beanstalk application when the customer load time surpasses your threshold.

Answer: C

Explanation:
Auto Scaling helps you ensure that you have the correct number of Amazon CC2 instances available to handle the load for your application. You create collections of
CC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Auto Scaling ensures that your group
never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Auto Scaling ensures that yourgroup never goes above this size. If you specify the desired capacity, either when you create the group or at any time thereafter.
Auto Scaling ensures that yourgroup has this many
instances. If you specify scaling policies, then Auto Scaling can launch or terminate instances as demand on your application increases or decreases.
Option A and B are invalid because Autoscaling is required to solve the issue to ensure the application can handle high traffic loads.
Option D is invalid because there is no Autoscaling template.
For more information on Autoscaling, please refer to the below document link: from AWS
http://docs.aws.amazon.com/autoscaling/latest/userguide/Whatl sAutoScaling.html

NEW QUESTION 62
A company's web application will be migrated to AWS. The application is designed so that there is no server-side code required. As part of the migration, the company would like to improve the security of the application by adding HTTP response headers, following the Open Web Application Security Project (OWASP) secure headers recommendations.
How can this solution be implemented to meet the security requirements using best practices?

A. Use an Amazon S3 bucket configured for website hosting, then set up server access logging on the S3 bucket to track user activit
B. Then configure the static website hosting and execute a scheduled AWS Lambda function to verify, and if missing, add security headers to the metadata.
C. Use an Amazon S3 bucket configured for website hosting, then set up server access logging on the S3 bucket to track user activit
D. Configure the static website hosting to return the required security headers.
E. Use an Amazon S3 bucket configured for website hostin
F. Create an Amazon CloudFront distribution that refers to this S3 bucket, with the origin response event set to trigger a Lambda@Edge Node.js function to add in the security headers.
G. set an Amazon S3 bucket configured for website hostin
H. Create an Amazon CloudFront distribution that refers to this S3 bucke
I. Set "Cache Based on Selected Request Headers" to "Whitelist," and add the security headers into the whitelist.

Answer: C

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/adding-http-security-headers-using-lambdaedge

NEW QUESTION 64
An ecommerce company uses a large number of Amazon EBS backed Amazon EC2 instances. To decrease manual work across all the instances, a DevOps engineer is tasked with automating restart actions when EC2 instance retirement events are scheduled.
How can this be accomplished?

A. Create a scheduled Amazon CloudWatch Events rule to execute an AWS Systems Manager automation document that checks if any EC2 instances are scheduled for retirement once a wee
B. If the instance is scheduled for retirement, the automation document will hibernate the instance.
C. Enable EC2 Auto Recovery on all of the instance
D. Create an AWS Config rule to limit the recovery to occur during a maintenance window only.
E. Reboot all EC2 instances during an approved maintenance window that is oil^ide of standard business hour
F. Set up Amazon CloudWatch alarms to send a notification in case any instance is failing EC2 instance status checks.
G. Set up an AWS Health Amazon CloudWatch Events rule to execute AWS Systems Manager automation documents that stop and start the EC2 instance when a retirement scheduled event occurs.

Answer: D

NEW QUESTION 66
A retail company wants to use AWS Elastic Beanstalk to host its online sales website running on Java. Since this will be the production website, the CTO has the following requirements for the deployment strategy:
*Zero downtime. While the deployment is ongoing, the current Amazon EC2 instances in service should remain in service. No deployment or any other action should be performed on the EC2 instances because they serve production traffic.
*A new fleet of instances should be provisioned for deploying the new application version.
*Once the new application version is deployed successfully in the new fleet of instances, the new instances should be placed in service and the old ones should be removed.
*The rollback should be as easy as possible. If the new fleet of instances fail to deploy the new application version, they should be terminated and the current instances should continue serving traffic as normal.
*The resources within the environment (EC2 Auto Scaling group, Elastic Load Balancing, Elastic Beanstalk DNS CNAME) should remain the same and no DNS change should be made.
Which deployment strategy will meet the requirements?

A. Use rolling deployments with a fixed amount of one instance at a time and set the healthy threshold to OK.
B. Use rolling deployments with additional batch with a fixed amount of one instance at a time and set the healthy threshold to OK.
C. launch a new environment and deploy the new application version there, then perform a CNAME swap between environments.
D. Use immutable environment updates to meet all the necessary requirements.

Answer: D

Explanation:
https://aws.amazon.com/about-aws/whats-new/2016/04/aws-elastic-beanstalk-adds-two-new-deployment-policie
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environmentmgmt-updates-immutable.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/command-options-general.html#command-options-gener

NEW QUESTION 71
A DevOps Engineer is building a multi-stage pipeline with AWS CodePipeline to build, verify, stage, test, and deploy an application. There is a manual approval stage required between the test and deploy stages. The development team uses a team chat tool with webhook support.
How can the Engineer configure status updates for pipeline activity and approval requests to post to the chat tool?

A. Create an AWS CloudWatch Logs subscription that filters on "detail-type": "CodePipeline PipelineExecution State Change." Forward that to an Amazon SNS topi
B. Add the chat webhook URL to the SNS topic as a subscriber and complete the subscription validation.
C. Create an AWS Lambda function that is triggered by the updating of AWS CloudTrail event
D. When a "CodePipeline Pipeline Execution State Change" event is detected in the updated events, send the event details to the chat webhook URL.
E. Create an AWS CloudWatch Events rule that filters on "CodePipeline Pipeline Execution State Change." Forward that to an Amazon SNS topi
F. Subscribe an AWS Lambda function to the Amazon SNS topic and have it forward the event to the chat webhook URL.
G. Modify the pipeline code to send event details to the chat webhook URL at the end of each stage.Parametrize the URL so each pipeline can send to a different URL based on the pipeline environment.

Answer: C

NEW QUESTION 73
A company uses AWS KMS with CMKs 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 SNS topic when keys are more than 90 days old.
B. Configure an Amazon CloudWatch Events event to launch an AWS Lambda function to call the AWS Trusted Advisor API and publish to an Amazon SNS topic
C. Develop an AWS Config custom rule that publishes to an Amazon SNS topic when keys are more than 90 days old
D. Configure AWS Security Hub to publish to an Amazon SNS topic when keys are more than 90 days old.

Answer: C

NEW QUESTION 77
A company has developed a Node.js web application which provides REST services to store and retrieve time series data. The web application is built by the Development team on company laptops, tested locally, and manually deployed to a single on-premises server, which accesses a local MySQL database. The company is starting a trial in two weeks, during which the application will undergo frequent updates based on customer feedback. The following requirements must be met:
*The team must be able to reliably build, test, and deploy new updates on a daily basis, without downtime or degraded performance.
*The application must be able to scale to meet an unpredictable number of concurrent users during the trial. Which action will allow the team to quickly meet these objectives?

A. Create two Amazon Lightsail virtual private servers for Node.js; one for test and one for production.Build the Node.js application using existing process and upload it to the new Lightsail test server using the AWS CL
B. Test the application, and if it passes all tests, upload it to the production serve
C. During the trial, monitor the production server usage, and if needed, increase performance by upgrading the instance type.
D. Develop an AWS CloudFormation template to create an Application Load Balancer and two Amazon EC2 instances with Amazon EBS (SSD) volumes in an Auto Scaling group with rolling updates enable
E. Use AWS CodeBuild to build and test the Node.js application and store it in an Amazon S3 bucke
F. Use user- data scripts to install the application and the MySQL database on each EC2 instanc
G. Update the stack to deploy new application versions.
H. Configure AWS Elastic Beanstalk to automatically build the application using AWS CodeBuild and to deploy it to a test environment that is configured to support auto scalin
I. Create a second Elastic Beanstalk environment for productio
J. Use Amazon RDS to store dat
K. When new versions of the applications have passed all tests, use Elastic Beanstalk "˜swap cname' to promote the test environment to production.
L. Modify the application to use Amazon DynamoDB instead of a local MySQL databas
M. Use AWS OpsWorks to create a stack for the application with a DynamoDB layer, an Application Load Balancer layer, and an Amazon EC2 instance laye
N. Use a Chef recipe to build the application and a Chef recipe to deploy the application to the EC2 instance laye
O. Use custom health checks to run unit tests on each instance with rollback on failure.

Answer: C

NEW QUESTION 78
A company is using AWS CodeBuild, AWS CodeDeploy, and AWS CodePipeline to deploy applications automatically to an Amazon EC2 instance. A DevOps Engineer needs to perform a security assessment scan of the operating system on every application deployment to the environment.
How should this be automated?

A. Use Amazon CloudWatch Events to monitor for Auto Scaling event notifications of new instances and configure CloudWatch Events to trigger an Amazon Inspector scan.
B. Use Amazon CloudWatch Events to monitor for AWS CodeDeploy notifications of a successful code deployment and configure CloudWatch Events to trigger an Amazon Inspector scan.
C. Use Amazon CloudWatch Events to monitor for CodePipeline notifications of a successful code deployment and configure CloudWatch Events to trigger an AWS X-Ray scan.
D. Use Amazon Inspector as a CodePipeline task after the successful use of CodeDeploy to deploy the code to the systems.

Answer: A

NEW QUESTION 79
An e-commerce company is running a web application in an AWS Elastic Beanstalk environment. In recent months, the average load of the Amazon EC2 instances has been increased to handle more traffic.
The company would like to improve the scalability and resilience of the environment. The Development team has been asked to decouple long-running tasks from the environment if the tasks can be executed asynchronously. Examples of these tasks include confirmation emails when users are registered to the platform, and processing images or videos. Also, some of the periodic tasks that are currently running within the web server should be offloaded.
What is the most time-efficient and integrated way to achieve this?

A. Create an Amazon SQS queue and send the tasks that should be decoupled from the Elastic Beanstalk web server environment to the SQS queu
B. Create a fleet of EC2 instances under an Auto Scaling grou
C. Use an AMI that contains the application to process the asynchronous tasks, configure the application to listen for messages within the SQS queue, and create periodic tasks by placing those into the cron in the operating syste
D. Create an environment variable within the Elastic Beanstalk environment with a value pointing to the SQS queue endpoint.
E. Create a second Elastic Beanstalk worker tier environment and deploy the application to process the asynchronous tasks ther
F. Send the tasks that should be decoupled from the original Elastic Beanstalk web server environment to the auto-generated Amazon SQS queue by the Elastic Beanstalk worker environmen
G. Place a cron.yaml file within the root of the application source bundle for the worker environment periodic task
H. Use environment links to link the web server environment with the worker environment.
I. Create a second Elastic Beanstalk web server tier environment and deploy the application to process the asynchronous task
J. Send the tasks that should be decoupled from the original Elastic Beanstalk web server to the auto-generated Amazon SQS queue by the Elastic Beanstalk web

server tier environmen
K. Place a cron.yaml file within the root of the application source bundle for the second web server tier environment with the necessary periodic task
L. Use environment links to link both web server environments.
M. Create an Amazon SQS queue and send the tasks that should be decoupled from the Elastic Beanstalk web server environment to the SQS queu
N. Create a fleet of EC2 instances under an Auto Scaling grou
O. Install and configure the application to listen for messages within the SQS queue from UserData and create periodic tasks by placing those into the cron in the operating syste
P. Create an environment variable within the Elastic Beanstalk web server environment with a value pointing to the SQS queue endpoint.

Answer: C

NEW QUESTION 83
A DevOps engineer must ensure all IAM entity configurations across multiple AWS accounts in AWS Organizations are compliant with corporate IAM policies.
Which combination of steps will accomplish this? (Select TWO.)

A. Enable AWS Trusted Advisor in Organizations for all accounts to report on noncompliant IAM entities.
B. Configure an AWS Config aggregator in the Organizations master account for all accounts
C. Deploy AWS Config rules to the master account in Organizations that match corporate IAM policies.
D. Apply an SCP in Organizations to ensure compliance of IAM entities.
E. Deploy AWS Config rules to all accounts in Organizations that match the corporate IAM policies.

Answer: BE

NEW QUESTION 85
Which Auto Scaling process would be helpful when testing new instances before sending traffic to them, while still keeping them in your Auto Scaling Group?

A. Suspend the process AZ Rebalance
B. Suspend the process Health Check
C. Suspend the process Replace Unhealthy
D. Suspend the process AddToLoadBalancer

Answer: D

Explanation:
If you suspend Ad dTo Load Balancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddTo Load Balancer process. Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does not add the instances that were launched while this process was suspended. You must register those instances manually.
Option A is invalid because this just balances the number of CC2 instances in the group across the Availability Zones in the region
Option B is invalid because this just checks the health of the instances. Auto Scaling marks an instance as unhealthy if Amazon CC2 or Clastic Load Balancing tells
Auto Scaling that the instance is unhealthy.
Option C is invalid because this process just terminates instances that are marked as unhealthy and later creates new instances to replace them.
For more information on process suspension, please refer to the below document link: from AWS
http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html

NEW QUESTION 88
A company wants to use AWS Systems Manager documents to bootstrap physical laptops for developers. The bootstrap code is stored in GitHub. A DevOps engineer has already created a Systems Manager activation, installed the Systems Manager agent with the registration code, and installed an activation ID on all the laptops.
Which set of steps should be taken next?

A. Configure the Systems Manager document to use the AWS-RunShellScript command to copy the files from GitHub to Amazon S3, then use the aws-downloadContent plugin with a source Type of S3.
B. Configure the Systems Manager document to use the aws-configurePackage plugin with an install action and point to the Git repository.
C. Configure the Systems Manager document to use the aws-downloadContent plugin with a sourceType of GitHub and sourceInfo with the repository details.
D. Configure the Systems Manager document to use the aws:softwareInventory plugin and run the script from the Git repository.

Answer: D

NEW QUESTION 90
A media customer has several thousand amazon EC2 instances in an AWS account. The customer is using a Slack channel for team communications and important updates. A DevOps Engineer was told to send all AWS-scheduled EC2 maintenance notifications to the company Slack channel.
Which method should the Engineer use to implement this process in the LEAST amount of steps?

A. Integrate AWS Trusted Advisor with AWS Confi
B. Based on the AWS Config rules created, the AWS Config event can invoke an AWS Lambda function to send notifications to the Slack channel.
C. Integrate AWS Personal Health Dashboard with Amazon CloudWatch Event
D. Based on the CloudWatch Events created, the event can invoke an AWS Lambda function to send notifications to the Slack channel.
E. Integrate EC2 events with Amazon CloudWatch monitorin
F. Based on the CloudWatch Alarm created, the alarm can invoke an AWS Lambda function to send EC2 maintenance notifications to the Slack channel.
G. Integrate AWS Support with AWS CloudTrai
H. Based on the CloudTrail lookup event created, the event can invoke an AWS Lambda function to pass EC2 maintenance notifications to the Slack channel.

Answer: B

Explanation:
https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html

NEW QUESTION 91
The management team at a company with a large on-premises OpenStack environment wants to move non-production workloads to AWS. An AWS Direct Connect connection has been provisioned and configured to connect the environments. Due to contractual obligations, the production workloads must remain on-premises, and will be moved to AWS after the next contract negotiation. The company follows Center for Internet Security (CIS) standards for hardening images; this configuration was developed using the company's configuration management system.
Which solution will automatically create an identical image in the AWS environment without significant overhead?

A. Write an AWS CloudFormation template that will create an Amazon EC2 instanc
B. Use cloud-unit to install the configuration management agent, use cfn-wait to wait for configuration management to successfully apply, and use an AWS Lambda-backed custom resource to create the AMI.
C. Log in to the console, launch an Amazon EC2 instance, and install the configuration management agent.When changes are applied through the configuration management system, log in to the console and create a new AMI from the instance.
D. Create a new AWS OpsWorks layer and mirror the image hardening standard
E. Use this layer as the baseline for all AWS workloads.
F. When a change is made in the configuration management system, a job in Jenkins is triggered to use the VM Import command to create an Amazon EC2 instance in the Amazon VP
G. Use lifecycle hooks to launch an AWS Lambda function to create the AMI.

Answer: D

Explanation:
https://www.brad-x.com/2015/10/01/importing-an-openstack-vm-into-amazon-ec2/ https://aws.amazon.com/ec2/vm-import/

NEW QUESTION 95
A company has a single Developer writing code for an automated deployment pipeline. The Developer is storing source code in an Amazon S3 bucket for each project. The company wants to add more Developers to the team but is concerned about code conflicts and lost work. The company also wants to build a test environment to deploy newer versions of code for testing and allow Developers to automatically deploy to both environments when code is changed in the repository.
What is the MOST efficient way to meet these requirements?

A. Create an AWS CodeCommit repository for each project, use the master branch for production code, and create a testing branch for code deployed to testin
B. Use feature branches to develop new features and pull requests to merge code to testing and master branches.
C. Create another S3 bucket for each project for testing code, and use an AWS Lambda function to promote code changes between testing and production bucket
D. Enable versioning on all buckets to prevent code conflicts.
E. Create an AWS CodeCommit repository for each project, and use the master branch for production and test code with different deployment pipelines for each environmen
F. Use feature branches to develop new features.
G. Enable versioning and branching on each S3 bucket, use the master branch for production code, and create a testing branch for code deployed to testin
H. Have Developers use each branch for developing in each environment.

Answer: A

NEW QUESTION 96
A startup company is developing a web application on AWS. It plans to use Amazon RDS for persistence and deploy the application to Amazon EC2 with an Auto Scaling group. The company would also like to separate the environments for development, testing, and production.
What is the MOST secure and flexible approach to manage the application configuration?

A. Create a property file to include the configuration and the encrypted password
B. Check in the property file to the source repository, package the property file with the application, and deploy the applicatio
C. Create an environment tag for the EC2 instances and tag the instances respectivel
D. The application will extract the necessary property values based on the environment tag.
E. Create a property file for each environment to include the environment-specific configuration and an encrypted passwor
F. Check in the property files to the source repositor
G. During deployment, use only the environment-specific property file with the applicatio
H. The application will read the needed property values from the deployed property file.
I. Create a property file for each environment to include the environment-specific configuratio
J. Create a private Amazon S3 bucket and save the property files in the bucke
K. Save the passwords in the bucket with AWS KMS encryptio
L. During deployment, the application will read the needed property values from the environment-specific property file in the S3 bucket.
M. Create a property file for each environment to include the environment-specific configuratio
N. Create a private Amazon S3 bucket and save the property files in the bucke
O. Save the encrypted passwords in the AWS Systems Manager Parameter Stor
P. Create an environment tag for the EC2 instances and tag the instances respectivel
Q. The application will read the needed property values from theenvironment-specific property file in the S3 bucket and the parameter store.

Answer: D

NEW QUESTION 101
A company wants to use Amazon ECS to provide a Docker container runtime environment. For compliance reasons, all Amazon EBS volumes used in the ECS cluster must be encrypted. Rolling updates will be made to the cluster instances and the company wants the instances drained of all tasks before being terminated.
How can these requirements be met? (Select TWO.)

A. Modify the default ECS AMI user data to create a script that executes docker rm ""f {id} for all running container instance
B. Copy the script to the /etc/ init.d/rc.d directory and execute chconfig enabling the script to run during operating system shutdown.
C. Use AWS CodePipeline to build a pipeline that discovers the latest Amazon-provided ECS AMI, then copies the image to an encrypted AMI outputting the encrypted AMI I
D. Use the encrypted AMI ID when deploying the cluster.
E. Copy the default AWS CloudFormation template that ECS uses to deploy cluster instance
F. Modify the template resource EBS configuration setting to set '"˜Encrypted: True' and include the AWS KMS alias: '"˜aws/ebs' to encrypt the AMI.
G. Create an Auto Scaling lifecycle hook backed by an AWS Lambda function that uses the AWS SDK to mark a terminating instance as DRAININ

H. Prevent the lifecycle hook from completing until the running tasks on the instance are zero.
I. Create an IAM role that allows the action ECS::EncryptedImag
J. Configure the AWS CLI and a profile to use this rol
K. Start the cluster using the AWS CLI providing the --use-encrypted-image and --kms-key arguments to the create-cluster ECS command.

Answer: CD

NEW QUESTION 102
A company is developing a web application's infrastructure using AWS CloudFormation. The database engineering team maintains the database resources in a CloudFormation template, and the software development team maintains the web application resources in a separate CloudFormation template As the scope of the application grows, the software development team needs to use resources maintained by the database engineering team However, both teams have their own review and lifecycle management processes that they want to keep Both teams also require resource-level change-set reviews The software development team would like to deploy changes to this template using their CI/CD pipeline.
Which solution will meet these requirements?

A. Create a stack export from the database CloudFormation template and import those references into the web application CloudFormation template
B. Create a CloudFormation nested stack to make cross-stack resource references and parameters available in both stacks.
C. Create a CloudFormation stack set to make cross-stack resource references and parameters available in both stacks
D. Create input parameters in the web application CloudFormation template and pass resource names and IDs from the database stack.

Answer: A

NEW QUESTION 107
A company updated the AWS CloudFormation template tor a critical business application. The stack update process Tailed due to an error in me updated template, and CloudFormation automatically began the stack rollback process Later, a DevOps engineer found the application was still unavailable, and that the stack was in the UPDATE_ROLLBACK_FALED state
Which combination of actions will allow the stack rollback to complete successful/? (Select TWO)

A. Attach the AWSCloudFormationFulAccess IAM policy to the CloudFormation role
B. Automatically heal the stack resources using CloudFormation drift detection.
C. Issue a ContinueUpdateRolback command from the CloudFormation console or AWS CLI
D. Manually the resources to match the expectations of the stack.
E. Update the existing CloudFormation stack using the original template

Answer: AB

NEW QUESTION 109
A company wants to automatically re-create its infrastructure using AWS CloudFormation as part of the company's quality assurance (QA) pipeline. For each QA run, a new VPC must be created in a single account, resources must be deployed into the VPC, and tests must be run against this new infrastructure. The company policy states that all VPCs must be peered with a central management VPC to allow centralized logging. The company has existing CloudFormation templates to deploy its VPC and associated resources.
Which combination of steps will achieve the goal in a way that is automated and repeatable? (Choose two.)

A. Create an AWS Lambda function that is invoked by an Amazon CloudWatch Events rule when a CreateVpcPeeringConnection API call is mad
B. The Lambda function should check the source of the peering request, accepts the request, and update the route tables for the management VPC to allow traffic to go over the peering connection.
C. In the CloudFormation template:Invoke a custom resource to generate unique VPC CIDR ranges for the VPC and subnets.Create a peering connection to the management VPC.Update route tables to allow traffic to the management VPC.
D. In the CloudFormation template:Use the Fn::Cidr function to allocate an unused CIDR range for the VPC and subnets.Create a peering connection to the management VPC.Update route tables to allow traffic to the management VPC.
E. Modify the CloudFormation template to include a mappings object that includes a list of /16 CIDR ranges for each account where the stack will be deployed.
F. Use CloudFormation StackSets to deploy the VPC and associated resources to multiple AWS accounts using a custom resource to allocate unique CIDR range
G. Create peering connections from each VPC to the central management VPC and accept those connections in the management VPC.

Answer: BD

NEW QUESTION 113
A devops team uses AWS CloudFormation to build their infrastructure. The security team is concerned about sensitive parameters, such as passwords, being exposed.
Which combination of steps will enhance the security of AWS CloudFormation? (Select THREE.)

A. Create a secure string with AWS KMS and choose a KMS encryption ke
B. Reference the ARN of the secure string, and give AWS CloudFormation permission to the KMS key for decryption.
C. Create secrets using the AWS Secrets Manager AWS::SecretsManager::Secret resource typ
D. Reference the secret resource return attributes in resources that need a password, such as an Amazon RDS database.
E. Store sensitive static data as secure strings in the AWS Systems Manager Parameter Stor
F. Use dynamic references in the resources that need access to the data.
G. Store sensitive static data in the AWS Systems Manager Parameter Store as string
H. Reference the stored value using types of Systems Manager parameters.
I. Use AWS KMS to encrypt the CloudFormation template.
J. Use the CloudFormation NoEcho parameter property to mask the parameter value.

Answer: ABD

NEW QUESTION 116
An Information Security policy requires that all publicly accessible systems be patched with critical OS security patches within 24 hours of a patch release. All instances are tagged with the Patch Group key set to 0. Two new AWS Systems Manager patch baselines for Windows and Red Hat Enterprise Linux (RHEL) with zero-day delay for security patches of critical severity were created with an auto-approval rule. Patch Group 0 has been associated with the new patch baselines.

Which two steps will automate patch compliance and reporting? (Select TWO.)

A. Create an AWS Systems Manager Maintenance Window and add a target with Patch Group 0. Add a task that runs the AWS-InstallWindowsUpdates document with a daily schedule.
B. Create an AWS Systems Manager Maintenance Window with a daily schedule and add a target with Patch Group 0. Add a task that runs the AWS-RunPatchBaseline document with the Install action.
C. Create an AWS Systems Manager State Manager configuratio
D. Associate the AWS-RunPatchBaseline task with the configuration and add a target with Patch Group 0.
E. Create an AWS Systems Manager Maintenance Window and add a target with Patch Group 0. Add a task that runs the AWS-ApplyPatchBaseline document with a daily schedule.
F. Use the AWS Systems Manager Run Command to associate the AWS-ApplyPatchBaseline document with instances tagged with Patch Group 0.

Answer: AC

NEW QUESTION 119
A DevOps Engineer is working with an application deployed to 12 Amazon EC2 instances across 3 Availability Zones. New instances can be started from an AMI image. On a typical day, each EC2 instance has 30% utilization during business hours and 10% utilization after business hours. The CPU utilization has an immediate spike in the first few minutes of business hours. Other increases in CPU utilization rise gradually.
The Engineer has been asked to reduce costs while retaining the same or higher reliability. Which solution meets these requirements?

A. Create two Amazon CloudWatch Events rules with schedules before and after business hours begin and en
B. Create two AWS Lambda functions, one invoked by each rul
C. The first function should stop nine instances after business hours end, the second function should restart the nine instances before the business day begins.
D. Create an Amazon EC2 Auto Scaling group using the AMI image, with a scaling action based on the Auto Scaling group's CPU Utilization average with a target of 75%. Create a scheduled action for the group to adjust the minimum number of instances to three after business hours end and reset to six before business hours begin.
E. Create two Amazon CloudWatch Events rules with schedules before and after business hours begin and en
F. Create an AWS CloudFormation stack, which creates an EC2 Auto Scaling group, with a parameter for the number of instance
G. Invoke the stack from each rule, passing a parameter value of three in the morning, and six in the evening.
H. Create an EC2 Auto Scaling group using the AMI image, with a scaling action based on the Auto Scaling group's CPU Utilization average with a target of 75%. Create a scheduled action to terminate nine instances each evening after the close of business.

Answer: B

NEW QUESTION 122
A company using AWS CodeCommit for source control wants to automate its continuous integration and continuous deployment pipeline on AWS in its development environment. The company has three requirements:
* 1. There must be a legal and a security review of any code change to make sure sensitive information is not leaked through the source code.
* 2. Every change must go through unit testing.
* 3. Every change must go through a suite of functional testing to ensure functionality. In addition, the company has the following requirements for automation:
* 1. Code changes should automatically trigger the CI/CD pipellline.
* 2. Any failure in the pipeline should notify devops-admin@xyz.com.
* 3. There must be an approval to stage the assets to Amazon S3 after tests have been performed.
What should a DevOps Engineer do to meet all of these requirements while following CI/CD best practices?

A. Commit to the development branch and trigger AWS CodePipeline from the development branc
B. Make an individual stage in CodePipeline for security review, unit tests, functional tests, and manual approva
C. Use Amazon CloudWatch metrics to detect changes in pipeline stages and Amazon SES for emailing devops- admin@xyz.com.
D. Commit to mainline and trigge r AWS CodePipeline from mainlin
E. Make an individual stage in CodePipeline for security review, unit tests, functional tests, and manual approva
F. Use AWS CloudTrail logs to detect changes in pipeline stages and Amazon SNS for emailing devops-admin@xyz.com.
G. Commit to the development branch and trigger AWS CodePipeline from the development branc
H. Make an individual stage in CodePipeline for security review, unit tests, functional tests, and manual approva
I. Use Amazon CloudWatch Events to detect changes in pipeline stages and Amazon SNS for emailing devops- admin@xyz.com.
J. Commit to mainline and trigger AWS CodePipeline from mainlin
K. Make an individual stage in CodePipeline for security review, unit tests, functional tests, and manual approva
L. Use Amazon CloudWatch Events to detect changes in pipeline stages and Amazon SES for emailingdevops-admin@xyz.com.

Answer: C

NEW QUESTION 127
A software company wants to automate the build process for a project where the code is stored in GitHub. When the repository is updated, source code should be compiled, tested, and pushed to Amazon S3.
Which combination of steps would address these requirements? (Select THREE.)

A. Add a buildspec.yml file to the source code with build instructions.
B. Configure a GitHub webhook to trigger a build every time a code change is pushed to the repository.
C. Create an AWS CodeBuild project with GitHub as the source repository.
D. Create an AWS CodeDeploy application with the Amazon EC2/On-Premises compute platform.
E. Create an AWS OpsWorks deployment with the install dependencies command.
F. Provision an Amazon EC2 instance to perform the build.

Answer: ACD

NEW QUESTION 132
A DevOps Engineer uses Docker container technology to build an image-analysis application. The application often sees spikes in traffic. The Engineer must automatically scale the application in response to customer demand while maintaining cost effectiveness and minimizing any impact on availability.
What will allow the FASTEST response to spikes in traffic while fulfilling the other requirements?

A. Create an Amazon ECS cluster with the container instances in an Auto Scaling grou
B. Configure the ECS service to use Service Auto Scalin
C. Set up Amazon CloudWatch alarms to scale the ECS service and cluster.
D. Deploy containers on an AWS Elastic Beanstalk Multicontainer Docker environmen
E. Configure Elastic Beanstalk to automatically scale the environment based on Amazon CloudWatch metrics.
F. Create an Amazon ECS cluster using Spot instance
G. Configure the ECS service to use Service Auto Scalin
H. Set up Amazon CloudWatch alarms to scale the ECS service and cluster.
I. Deploy containers on Amazon EC2 instance
J. Deploy a container scheduler to schedule containers onto EC2 instance
K. Configure EC2 Auto Scaling for EC2 instances based on available Amazon CloudWatch metrics.

Answer: A

Explanation:
https://aws.amazon.com/blogs/compute/automatic-scaling-with-amazon-ecs/

NEW QUESTION 136
A development team manages website deployments using AWS CodeDeploy blue/green deployments. The application is running on Amazon EC2 instances behind an Application Load Balancer in an Auto Scaling group.
When deploying a new revision, the team notices the deployment eventually fails, but it takes a long time to fail. After further inspection, the team discovers the AllowTrafflc lifecycle event ran for an hour and eventually failed without providing any other information. The team wants to ensure failure notices are delivered more quickly while maintaining application availability even upon failure.
Which combination of actions should be taken to meet these requirements? (Select TWO.)

A. Change the deployment configuration to CodeDeployDefault.AIIAtOnce to speed up the deployment process by deploying to all of the instances at the same time.
B. Create a CodeDeploy trigger for the deployment failure event and make the deployment fail as soon as a single health check failure is detected.
C. Reduce the HealthCheckIntervalSeconds and UnhealthyThresholdCount values within the target group health checks to decrease the amount of time it takes for the application to be considered unhealthy.
D. Use the appspec.yml file to run a script on the AllowTraffic hook to perform lighter health checks on the application instead of making CodeDeploy wait for the target group health checks to pass.
E. Use the appspec.yml file to run a script on the BeforeAllowTraffic hook to perform health checks on the application and fail the deployment if the health checks performed by the script are not successful.

Answer: AE

NEW QUESTION 139
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 143
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 Cloud Formation. 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 in source control.
Which solution will accomplish this?

A. In the CloudFormation 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-init 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 init metadat

K. Place the configuration file content in the metadat
L. Configure the cfn-init script to run when the instance is launched, and configure thecfn-hup script to poll for updates to the configuration.

Answer: B

NEW QUESTION 144
A company wants to migrate its content sharing web application hosted on Amazon EC2 to a serverless architecture. The company currently deploys changes to its application by creating a new Auto Scaling group of EC2 instances and a new Elastic Load Balancer and then shifting the traffic away using an Amazon Route 53 weighted routing policy
For its new serverless application, the company is planning to use Amazon API Gateway and AWS Lambda The company will need to update its deployment processes to work with the new application. It will also need to retain the ability to test new features on a small number of users before rolling the features out to the entire user base
Which deployment strategy will meet these requirements?

A. Use AWS CDK to deploy API Gateway and Lambda function
B. When code needs to be changed, update the AWS CloudFormation stack and deploy the new version of the APIs and Lambda function
C. Use a Route 53 failover routing policy for the canary release strategy.
D. Use AWS CloudFormation to deploy API Gateway and Lambda functions using Lambda function versions When code needs to be changed, update the CloudFormation stack with the new Lambda code and update the API versions using a canary release strategy Promote the new version when testing is complete.
E. Use AWS Elastic Beanstalk to deploy API Gateway and Lambda functions When code needs to be changed, deploy a new version of the API and Lambda function
F. Shift traffic gradually using an Elastic Beanstalk blue/green deployment.
G. Use AWS OpsWorks to deploy API Gateway in the service layer and Lambda functions in a custom laye
H. When code needs to be changed, use OpsWorks to perform a blue/green deployment and shift traffic gradually

Answer: B

NEW QUESTION 149
A company is implementing AWS CodePipeline to automate its testing process. The company wants to be notified when the execution state fails and used the following custom event pattern in Amazon CloudWatch:

```
{
  "source": [
    "aws.codepipeline"
  ],
  "detail-type": [
    "CodePipline Action Execution State Change"
  ],
  "detail": {
    "state": [
      "FAILED"
    ],
    "type": {
      "category": ["Approval"]
    }
  }
}
```

Which type of events will match this event pattern?

A. Failed deploy and build actions across all the pipelines.
B. All rejected or failed approval actions across all the pipelines.
C. All the events across all pipelines.
D. Approval actions across all the pipelines.

Answer: B

Explanation:
https://docs.aws.amazon.com/codepipeline/latest/userguide/detect-state-changes-cloudwatch-events.html

NEW QUESTION 151
A company is beginning to move to the AWS Cloud. Internal customers are classified into two groups according to their AWS skills: beginners and experts.
The DevOps Engineer needs to build a solution to allow beginners to deploy a restricted set of AWS architecture blueprints expresses as AWS CloudFormation templates. Deployment should only be possible on predetermined Virtual Private Clouds (VPCs). However, expert users should be able to deploy blueprints without constraints. Experts should also be able to access other AWS services, as needed.
How can the Engineer implement a solution to meet these requirements with the LEAST amount of overhead?

A. Apply constraints to the parameters in the templates, limiting the VPCs available for deployment
B. Store the templates on Amazon S3. Create an IAM group for beginners and give them access to the templates and CloudFormatio

C. Create a separate group for experts, giving them access to the templates, CloudFormation, and other AWS services.
D. Store the templates on Amazon S3. Use AWS Service Catalog to create a portfolio of products based on those template
E. Apply template constraints to the products with rules limiting VPCs available for deployment
F. Create an IAM group for beginners giving them access to the portfoli
G. Create a separate group for experts giving them access to the templates, CloudFormation, and other AWS services.
H. Store the templates on Amazon S3. Use AWS Service Catalog to create a portfolio of products based on those template
I. Create an IAM role restricting VPCs available for creation of AWS resource
J. Apply alaunch constraint to the products using this rol
K. Create an IAM group for beginners giving them access to the portfoli
L. Create a separate group for experts giving them access to the portfolio and other AWS services.
M. Create two templates for each architecture blueprint where only one of them limits the VPC available for deployment
N. Store the templates in Amazon DynamoD
O. Create an IAM group for beginners giving them access to the constrained templates and CloudFormatio
P. Create a separate group for experts giving them access to the unconstrained templates, CloudFormation, and other AWS services.

Answer: B

NEW QUESTION 152
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 on
B. Configure API Gateway to use a canary release deployment to send a small subset of user traffic to the new environment.
C. Introduce changes as a separate environment parallel to the existing on
D. Update the application's DNS alias records to point to the new environment.
E. Introduce changes as a separate target group behind the existing Application Load Balance
F. Configure API Gateway to route user traffic to the new target group in steps.
G. Introduce changes as a separate target group behind the existing Application Load Balance
H. Configure API Gateway to route all traffic to the Application Load Balancer, which then sends the traffic to the new target group.

Answer: A

NEW QUESTION 155
According to Information Security Policy, changes to the contents of objects inside production Amazon S3 bucket that contain encrypted secrets should only be made by a trusted group of administrators.
How should a DevOps Engineer create real-time, automated checks to meet this requirement?

A. Create an AWS Lambda function that is triggered by Amazon S3 data events for object changes and that also checks the IAM user's membership in an administrator's IAM role.
B. Create a periodic AWS Config rule to query Amazon S3 Logs for changes and to check the IAM user's membership in an administrator's IAM role.
C. Create a metrics filter for Amazon CloudWatch logs to check for Amazon S3 bucket-level permission changes and to check the IAM user's membership in an administrator's IAM role.
D. Create a periodic AWS Config rule to query AWS CloudTrail logs for changes to the Amazon S3 bucket-level permissions and to check the IAM user's membership in an administrator's IAM role.

Answer: A

NEW QUESTION 160
A company wants to implement a CI/CD pipeline for building and testing its mobile apps. A DevOps Engineer has been given the following requirements: Use AWS CodePipeline to orchestrate the workflow. Test the application on real devices. Trigger a notification. Stage the application binary on a production bucket in a different account. Make the application binary publicly accessible. Which sequence of actions should the Engineer perform in the pipeline to meet the requirements?

A. Use AWS CodeCommit as the code source and AWS CodeDeploy to compile and package the applicatio
B. Use CodeDeploy to deploy the application binary to an AWS Lambda function for testin
C. Use a third-party library on AWS Lambda to simulate the device platfor
D. Allow a Lambda role to upload to the production Amazon S3 bucke
E. Make the binary publicly accessibl
F. Trigger notifications using Amazon SNS.
G. Use GitHub as the code source and AWS Lambda to compile and package the applicatio
H. Use another Lambda function to run unit tests and deliver the application binary to a development bucke
I. Use the binary from the development bucket and install the application on a personal device for testin
J. Deliver the binary to the production bucket after approva
K. Trigger notifications using Amazon SNS.
L. Use an Amazon S3 bucket as the code source and AWS CodeBuild to compile and package the applicatio
M. Use AWS CodeDeploy to deploy the application binary to a device farm for testin
N. Deliver the binary to the production S3 bucke
O. Use an S3 bucket policy to allow public read on the productionS3 bucke
P. Trigger notifications using an Amazon CloudWatch Events rule with Amazon SNS.
Q. Use AWS CodeCommit as the code source and AWS CodeBuild to compile and package the applicatio
R. Invoke an AWS Lambda function that uploads the application binary to a device farm for testin
S. Deliver the binary to the production Amazon S3 bucke
T. Use an S3 bucket policy to allow public read on the production S3 bucke
. Trigger notifications by using an Amazon CloudWatch Events rule.

Answer: D

NEW QUESTION 163
A DevOps Engineer is responsible for the deployment of a PHP application. The Engineer is working in a hybrid deployment, with the application running on both on-premises servers and Amazon EC2 instances. The application needs access to a database containing highly confidential information. Application instances need access to database credentials, which must be encrypted at rest and in transit before reaching the instances.
How should the Engineer automate the deployment process while also meeting the security requirements?

A. Use AWS Elastic Beanstalk with a PHP platform configuration to deploy application packages to theinstance
B. Store database credentials on AWS Systems Manager Parameter Store using the Secure String data typ
C. Define an IAM role for Amazon EC2 allowing access, and decrypt only the database credential
D. Associate this role to all the instances.
E. Use AWS CodeDeploy to deploy application packages to the instance
F. Store database credentials on AWS Systems Manager Parameter Store using the Secure String data typ
G. Define an IAM policy for allowing access, and decrypt only the database credential
H. Attach the IAM policy to the role associated to the instance profile for CodeDeploy-managed instances, and to the role used for on-premises instances registration on CodeDeploy.
I. Use AWS CodeDeploy to deploy application packages to the instance
J. Store database credentials on AWS Systems Manager Parameter Store using the Secure String data typ
K. Define an IAM role with an attached policy that allows decryption of the database credential
L. Associate this role to all the instances and on-premises servers.
M. Use AWS CodeDeploy to deploy application packages to the instance
N. Store database credentials in the AppSpec fil
O. Define an IAM policy for allowing access to only the database credential
P. Attach the IAM policy to the role associated to the instance profile for CodeDeploy-managed instances and the role used for on-premises instances registration on CodeDeploy

Answer: B

NEW QUESTION 164
A defect was discovered in production and a new sprint item has been created for deploying a hotfix. However, any code change must go through the following steps before going into production:
*Scan the code for security breaches, such as password and access key leaks. Run the code through extensive, long running unit tests.
Which source control strategy should a DevOps Engineer use in combination with AWS CodePipeline to complete this process?

A. Create a hotfix tag on the last commit of the master branc
B. Trigger the development pipeline from the hotfix ta
C. Use AWS CodeDeploy with Amazon ECS to do a content scan and run unit test
D. Add a manual approval stage that merges the hotfix tag into the master branch.
E. Create a hotfix branch from the master branc
F. Triger the development pipeline from the hotfix branch.Use AWS CodeBuild to do a content scan and run unit test
G. Add a manual approval stage that merges the hotfix branch into the master branch.
H. Create a hotfix branch from the master branc
I. Triger the development pipeline from the hotfix branch.Use AWS Lambda to do a content scan and run unit test
J. Add a manual approval stage that merges the hotfix branch into the master branch.
K. Create a hotfix branch from the master branc
L. Create a separate source stage for the hotfix branch in the production pipelin
M. Trigger the pipeline from the hotfix branc
N. Use AWS Lambda to do a content scan and use AWS CodeBuild to run unit test
O. Add a manual approval stage that merges the hotfix branch into the master branch.

Answer: B

NEW QUESTION 168
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 using reserved concurrency. The Lambda function processes order messages from an 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 will diagnose and resolve the delay? (Select TWO.)

A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and increase the Lambda function concurrency limit.
B. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and configure a redrive policy on the SQS queue.
C. Check the NumberOfMessagesSent metric for the SQS queue and increase the SQS queue visibility timeout.
D. Check the ThrottledWriteRequests metric for the DynamoDB table and increase the maximum write capacity units for the table's Auto Scaling policy.
E. Check the Throttles metric for the Lambda function and increase the Lambda function timeout.

Answer: AB

NEW QUESTION 173
A company has a mission-critical application on AWS that uses automatic scaling. The company wants the deployment lifecycle to meet the following parameters
•The application must be deployed one instance at a time to ensure the remaining fleet continues to serve traffic.
•the application is CPU intensive and must ho closely monitored
•the deployment must automatically roll back if the CPU utilization of the deployment instance exceeds 85% Which solution will meet these requirements'?

A. Use AWS CloudForrnation to create an AWS Step Functions state machine and Auto Scaling lifecycle hooks to move to one instance at a time into a wait stat
B. Use AWS Systems Manager automation to deploy the update to each instance and move it back into the Auto Scaling group using the heartbeat timeout
C. Use AWS CodeDeploy with Amazon EC2 Auto Scaling Configure an alarm tied to the CPU utilization metric Use the CodeDeployDefault OneAtAtime configuration as a deployment strategy Configure automatic rollbacks within the deployment group to roll back the deployment if the alarm thresholds are breached
D. Use AWS Elastic Beanstalk for load balancing and AWS Auto Scaling Configure an alarm tied to the CPU utilization metric Configure rolling deployments with a fixed batch size of one instance Enable enhanced health to monitor the status of the deployment and roll back based on the alarm previously created
E. Use AWS Systems Manager to perform a blue/green deployment with Amazon EC2 Auto Scaling Configure an alarm tied to the CPU utilization metric Deploy updates one at a time Configure automatic rollbacks within the Auto Scaling group to roll back the deployment if the alarm thresholds are breached.

Answer: B

NEW QUESTION 174
A company has a hybrid architecture solution in which some legacy systems remain on-premises, while a specific cluster of servers is moved to AWS. The company cannot reconfigure the legacy systems, so the cluster nodes must have a fixed hostname and local IP address for each server that is part of the cluster. The DevOps Engineer must automate the configuration for a six-node cluster with high availability across three Availability Zones (AZs), placing two elastic network interfaces in a specific subnet for each AZ. Each node's hostname and local IP address should remain the same between reboots or instance failures. Which solution involves the LEAST amount of effort to automate this task?

A. Create an AWS Elastic Beanstalk application and a specific environment for each server of the cluster.For each environment, give the hostname, elastic network interface, and AZ as input parameter
B. Use the local health agent to name the instance and attach a specific elastic network interface based on the current environment.
C. Create a reusable AWS CloudFormation template to manage an Amazon EC2 Auto Scaling group with a minimum size of 1 and a maximum size of 1. Give the hostname, elastic network interface, and AZ as stack parameter
D. Use those parameters to set up an EC2 instance with EC2 Auto Scaling and a user datascript to attach to the specific elastic network interfac
E. Use CloudFormation nested stacks to nest the template six times for a total of six nodes needed for the cluster, and deploy using the master template.
F. Create an Amazon DynamoDB table with the list of hostnames subnets, and elastic network interfaces to be use
G. Create a single AWS CloudFormation template to manage an Auto Scaling group with a minimum size of 6 and a maximum size of 6. Create a programmatic solution that is installed in each instance that will lock/release the assignment of each hostname and local IP address, depending on the subnet in which a new instance will be launched.
H. Create a reusable AWS CLI script to launch each instance individually, which will name the instance, place it in a specific AZ, and attach a specific elastic network interfac
I. Monitor the instances and in the event of failure, replace the missing instance manually by running the script again.

Answer: B

NEW QUESTION 178
A DevOps Engineer has been asked by the Security team to ensure that AWS CloudTrail files are not tampered with after being created. Currently, there is a process with multiple trails, using AWS IAM to restrict access to specific trails. The Security team wants to ensure they can trace the integrity of each file and make sure there has been no tampering.
Which option will require the LEAST effort to implement and ensure the legitimacy of the file while allowing the Security team to prove the authenticity of the logs?

A. Create an Amazon CloudWatch Events rule that triggers an AWS Lambda function when a new file is delivere
B. Configure the Lambda function to perform an MD5 hash check on the file, store the name and location of the file, and post the returned hash to an Amazon DynamoDB tabl
C. The Security team can use the values stored in DynamoDB to verify the file authenticity.
D. Enable the CloudTrail file integrity feature on an Amazon S3 bucke
E. Create an IAM policy that grants the Security team access to the file integrity logs stored in the S3 bucket.
F. Enable the CloudTrail file integrity feature on the trai
G. Use the digest file created by CloudTrail to verify the integrity of the delivered CloudTrail files.
H. Create an AWS Lambda function that is triggered each time a new file is delivered to the CloudTrail bucke
I. Configure the Lambda function to execute an MD5 hash check on the file, and store the result on a tag in an Amazon S3 objec
J. The Security team can use the information on the tag to verify the integrity of the file.

Answer: C

Explanation:
https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html

NEW QUESTION 181
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with Our Prep Materials Via below:

https://www.certleader.com/AWS-Certified-DevOps-Engineer-Professional-dumps.html