Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1
A DevOps Engineer wants to prevent Developers from pushing updates directly to the company's master branch in AWS CodeCommit. These updates should be approved before they are merged.
Which solution will meet these requirements?

A. Configure an IAM role for the Developers with access to CodeCommit and an explicit deny for write actions when the reference is the maste
B. Allow Developers to use feature branches and create a pull request when a feature is complet
C. Allow an approver to use CodeCommit to view the changes and approve the pull requests.
D. Configure an IAM role for the Developers to use feature branches and create a pull request when a feature is complet
E. Allow CodeCommit to test all code in the feature branches, and dynamically modify the IAM role to allow merging the feature branches into the maste
F. Allow an approver to use CodeCommit to view the changes and approve the pull requests.
G. Configure an IAM role for the Developers to use feature branches and create a pull request when a feature is complet
H. Allow CodeCommit to test all code in the feature branches, and issue a new AWS Security Token Service (STS) token allowing a one-time API call to merge the feature branches into the maste
I. Allow an approver to use CodeCommit to view the changes and approve the pull requests.
J. Configure an IAM role for the Developers with access to CodeCommit and attach an access policy to the CodeCommit repository that denies the Developers role access when the reference is maste
K. Allow Developers to use feature branches and create a pull request when a feature is complet
L. Allow an approver to use CodeCommit to view the changes and approve the pull requests.

Answer: D

NEW QUESTION 2
A company has migrated its container-based applications to Amazon EKS and want to establish automated email notifications. The notifications sent to each email address are for specific activities related to EXS components. The solution will include Amazon SNS topics and an AWS Lambda function to evaluate incoming log events and publish messages to the correct SNS topic.
Which logging solution will support these requirements?

A. Enable Amazon CloudWatch Logs to log the EKS component
B. Create a CloudWatch subscription filterfor each component with Lambda as the subscription feed destination.
C. Enable Amazon CloudWatch Logs to log the EKS component
D. Create CloudWatch Logs Insights queries linked to Amazon CloudWatch Events events that trigger Lambda.
E. Enable Amazon S3 logging for the EKS component
F. Configure an Amazon CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
G. Enable Amazon S3 logging for the EKS component
H. Configure S3 PUT Object event notifications with AWS Lambda as the destination.

Answer: A

NEW QUESTION 3
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 4
A company is using AWS CodeDeploy to automate software deployment. The deployment must meet these requirements:
*A number of instances must be available to serve traffic during the deployment. Traffic must be balanced across those instances, and the instances must automatically heal in the event of failure.
*A new fleet of instances must be launched for deploying a new revision automatically, with no manual provisioning.
*Traffic must be rerouted to the new environment to half of the new instances at a time. The deployment should succeed if traffic is rerouted to at least half of the instances; otherwise, it should fail.
*Before routing traffic to the new fleet of instances, the temporary files generated during the deployment process must be deleted.
*At the end of a successful deployment, the original instances in the deployment group must be deleted immediately to reduce costs.
How can a DevOps Engineer meet these requirements?

A. Use an Application Load Balancer and an in-place deploymen
B. Associate the Auto Scaling group with the deployment grou
C. Use the Automatically copy option, and use CodeDeployDefault.OneAtAtime as the deployment configuratio
D. Instruct AWS CodeDeploy to terminate the original Auto Scaling group instances in the deployment group, and use the AllowTraffic hook within appspec.yml to delete the temporary files.
E. Use an Application Load Balancer and a blue/green deploymen
F. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment grou
G. Use the Automatically copy Auto Scaling group option, create a custom deployment configuration with minimum healthy hosts defined as 50%, and assign the

configuration to the deployment grou
H. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeBlock Traffic hook within appsec.yml to delete the temporary files.
I. Use an Application Load Balancer and a blue/green deploymen
J. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment grou
K. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault HalfAtAtime as the deployment configuratio
L. Instruct AWS CodeDeploy to terminate the original isntances in the deployment group, and use the BeforeAllowTraffic hook within appspec.yml to delete the temporary files.
M. Use an Application Load Balancer and an in-place deploymen
N. Associate the Auto Scaling group and Application Load Balancer target group with the deployment grou
O. Use the Automatically copy AutoScaling group option, and use CodeDeployDefault AllatOnce as a deployment configuratio
P. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BlockTraffic hook within appsec.yml to delete the temporary files.

Answer: C

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html
https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueGreenDeploymentConfiguration.html

NEW QUESTION 5
A company wants 10 use AWS development tools to replace Its current bash deployment scripts. The company currently deploys a LAMP application to a group of Amazon EC2 instances behind an Application Load Balancer (ALB). During the deployments, the company unit tests the committed application, stops and starts services, unregisters and re-registers instances with the load balancer, and updates Me permissions. The company wants to maintain the same deployment functionality through the shift to using AWS services.
Which solution will meet these requirements?

A. Use AWS CodeBuild to test the applicatio
B. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services, and deregister and register instances with the ALB Use the appspec.yml file to update file permissions without a custom script.
C. Use AWS CodePipeline to move the application from the AWS CodeCommit repository to AWS CodeDeplo
D. Use CodeDeploy's deployment group to test the application, unregister and reregister instances with the AL
E. and restart service
F. Use the appspec.yml file to update file permissions without a custom script.
G. Use AWS CodePipeline to move the application source code from the AWS CodeCommit repository to AWS CodeDeplo
H. Use CodeDeploy to test the applicatio
I. Use CodeDeploy's appspec.yml file to restart services and update permissions without a custom scrip
J. Use AWS CodeBuild to unregister and re-register instances with the ALB.
K. Use AWS CodePipeline to trigger AWS CodeBuild to test the application Use bash scripts invoked by AWS CodeDeploy's appspec yml file to restart service
L. Unregister and re-register theinstances in the AWS CodeDeploy deployment group with the AL
M. Update the appspec.yml file to update file permissions without a custom script.

Answer: D

NEW QUESTION 6
A DevOps Engineer has a single Amazon DynamoDB table that received shipping orders and tracks inventory. The Engineer has three AWS Lambda functions reading from a DymamoDB stream on that table. The Lambda functions perform various functions such as doing an item count, moving items to Amazon Kinesis Data Firehose, monitoring inventory levels, and creating vendor orders when parts are low.
While reviewing logs, the Engineer notices the Lambda functions occasionally fail under increased load, receiving a stream throttling error.
Which is the MOST cost-effective solution that requires the LEAST amount of operational management?

A. Use AWS Glue integration to ingest the DynamoDB stream, then migrate the Lambda code to an AWS Fargate task.
B. Use Amazon Kinesis streams instead of DynamoDB streams, then use Kinesis analytics to trigger the Lambda functions.
C. Create a fourth Lambda function and configure it to be the only Lambda reading from the strea
D. Then use this Lambda function to pass the payload to the other three Lambda functions.
E. Have the Lambda functions query the table directly and disable DynamoDB stream
F. Then have the Lambda functions query from a global secondary index.

Answer: C

NEW QUESTION 7
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 8
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

A. Set up Amazon Route 53 health checks to fail over from any Amazon EC2 instance that is currently being deployed to.
B. Using AWS CloudFormation, create a test stack for validating the code, and then deploy the code to each production Amazon EC2 instance.
C. Create a new load balancer with new Amazon EC2 instances, carry out the deployment, and then switch DNS over to the new load balancer using Amazon Route 53 after testing.
D. Launch more Amazon EC2 instances to ensure high availability, de-register each Amazon EC2 instance from the load balancer, upgrade it, and test it, and then register it again with the load balancer.

Answer: C

Explanation:
The below diagram shows how this can be done C:\Users\wk\Desktop\mudassar\Untitled.jpg

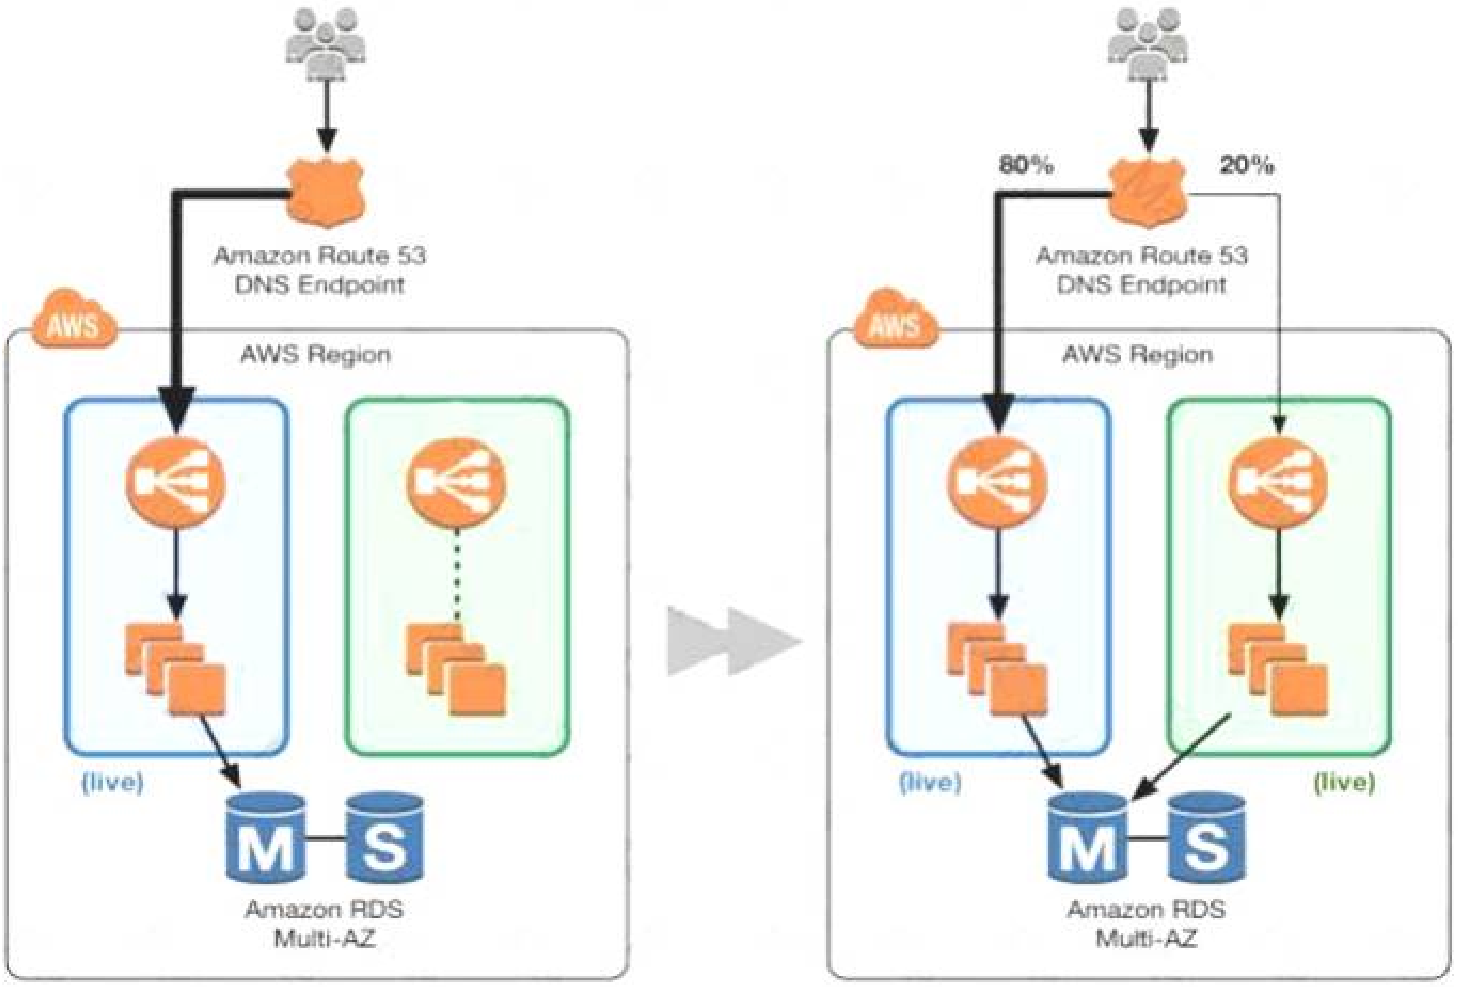

1) First create a new ELB which will be used to point to the new production changes.
2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80-20% traffic scenario. This is the normal case, the % can be changed based on the requirement.
3) Finally when all changes have been tested, Route53 can be set to 100% for the new ELB.
Option A is incorrect because this is a failover scenario and cannot be used for Blue green deployments. In Blue Green deployments, you need to have 2 environments running side by side.
Option B is incorrect, because you need to a have a production stack with the changes which will run side by side.
Option D is incorrect because this is not a blue green deployment scenario. You cannot control which users will go the new EC2 instances.
For more information on blue green deployments, please refer to the below document link: from AWS
https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 9
A DevOps engineer is troubleshooting deployments to a new application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. Instances sometimes come online before they are ready, which is leading to increased error rates among users. The current health check configuration gives instances a 60-second grace period and considers instances healthy after two 200 response codes from /index.php, a page that may respond intermittently during the deployment process. The development team wants instances to come online as soon as possible.
Which strategy would address this issue?

A. Increase the instance grace period from 60 seconds to 180 seconds, and the consecutive health check requirement from 2 to 3.
B. Increase the instance grace period from 60 seconds to 120 seconds, and change the response code requirement from 200 to 204.
C. Modify the deployment script to create a /health-check.php file when the deployment begins, then modify the health check path to point to that file.
D. Modify the deployment script to create a /health-check.php file when all tasks are complete, then modify the health check path to point to that file.

Answer: D

NEW QUESTION 10

A mobile application running on eight Amazon EC2 instances is relying on a third-party API endpoint. The thirdparty service has a high failure rate because of limited capacity, which is expected to be resolved in a few weeks. In the meantime, the mobile application developers have added a retry mechanism and are logging failed API requests. A DevOps Engineer must automate the monitoring of application logs and count the specific error messages; if there are more than 10 errors within a 1-minute window, the system must issue an alert. How can the requirements be met with MINIMAL management overhead?

A. Install the Amazon CloudWatch Logs agent on all instances to push the application logs to CloudWatch Log
B. Use metric filters to count the error messages every minute, and trigger a CloudWatch alarm if the count exceeds 10 errors.
C. Install the Amazon CloudWatch Logs agent on all instances to push the access logs to CloudWatch Log
D. Create CloudWatch Events rule to count the error messages every minute, and trigger a CloudWatch alarm if the count exceeds 10 errors.
E. Install the Amazon CloudWatch Logs agent on all instances to push the application logs to CloudWatchLog
F. Use a metric filter to generate a custom CloudWatch metric that records the number of failures and triggers a CloudWatch alarm if the custom metric reaches 10 errors in a 1-minute period.
G. Deploy a custom script on all instances to check application logs regularly in a cron jo
H. Count the number of error messages every minute, and push a data point to a custo
I. CloudWatch metri
J. Trigger a CloudWatch alarm if the custom metric reaches 10 errors in a 1-minute period.

Answer: C

NEW QUESTION 10
A company has deployed several applications globally. Recently, Security Auditors found that few Amazon EC2 instances were launched without Amazon EBS disk encryption. The Auditors have requested a report detailing all EBS volumes that were not encrypted in multiple AWS accounts and regions. They also want to be notified whenever this occurs in future.
How can this be automated with the LEAST amount of operational overhead?

A. Create an AWS Lambda function to set up an AWS Config rule on all the target account
B. Use AWS Config aggregators to collect data from multiple accounts and region
C. Export the aggregated report to an Amazon S3 bucket and use Amazon SNS to deliver the notifications.
D. Set up AWS CloudTrail to deliver all events to an Amazon S3 bucket in a centralized accoun
E. Use the S3 event notification feature to invoke an AWS Lambda function to parse AWS CloudTrail logs whenever logs are delivered to the S3 bucke
F. Publish the output to an Amazon SNS topic using the same Lambda function.
G. Create an AWS CloudFormation template that adds an AWS Config managed rule for EBS encryption.Use a CloudFormation stack set to deploy the template across all accounts and region
H. Store consolidated evaluation results from config rules in Amazon S3. Send a notification using Amazon SNS when non- compliant resources are detected.
I. Using AWS CLI, run a script periodically that invokes the aws ec2 describe-volumes query with a JMESPATH query filte
J. Then, write the output to an Amazon S3 bucke
K. Set up an S3 event notification to send events using Amazon SNS when new data is written to the S3 bucket.

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/aws-config-update-aggregate-compliance-data-across-accounts-regions/
https://docs.aws.amazon.com/config/latest/developerguide/aws-config-managed-rules-cloudformation-templates

NEW QUESTION 12
A company has containerized all of its in-house quality control applications. The company is running Jenkins on Amazon EC2. which requires patching and upgrading. The compliance officer has requested a DevOps engineer begin encrypting build artifacts since they contain company intellectual property.
What should the DevOps engineer do to accomplish this in the MOST maintainable manner?

A. Automate patching and upgrading using AWS Systems Manager on EC2 instances and encrypt Amazon EBS volumes by default.
B. Deploy Jenkins to an Amazon ECS cluster and copy build artifacts to an Amazon S3 bucket with defaultencryption enabled.
C. Leverage AWS CodePipeline with a build action and encrypt the artifacts using AWS Secrets Manager.
D. Use AWS CodeBuild with artifact encryption to replace the Jenkins instance running on Amazon EC2.

Answer: D

NEW QUESTION 13
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 16
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 18
A company indexes all of its Amazon CloudWatch Logs on Amazon ES and uses Kibana to view a dashboard for actionable insight. The company wants to restrict user access to Kibana by user
Which actions can a DevOps Engineer take to meet this requirement? (Select TWO.)

A. Create a proxy server with user authentication in an Auto Scaling group and restrict access of the Amazon ES endpoint to an Auto Scaling group tag
B. Create a proxy server with user authentication and an Elastic IP address and restrict access of the Amazon ES endpoint to the IP address
C. Create a proxy server with AWS IAM user and restrict access of the Amazon ES endpoint to the IAM user
D. Use AWS SSO to offer user name and password protection for Kibana
E. Use Amazon Cognito to offer user name and password protection for Kibana

Answer: BE

NEW QUESTION 22
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 25
A company wants to ensure that their EC2 instances are secure. They want to be notified if any new vulnerabilities are discovered on their instances, and they also want an audit trail of all login activities on the instances.
Which solution will meet these requirements?

A. Use AWS Systems Manager to detect vulnerabilities on the EC2 instance
B. Install the Amazon Kinesis Agent to capture system logs and deliver them to Amazon S3.
C. Use AWS Systems Manager to detect vulnerabilities on the EC2 instance
D. Install the Systems Manager Agent to capture system logs and view login activity in the CloudTrail console.
E. Configure Amazon CloudWatch to detect vulnerabilities on the EC2 instance
F. Install the AWS Config daemon to capture system logs and view them in the AWS Config console.
G. Configure Amazon Inspector to detect vulnerabilities on the EC2 instance
H. Install the Amazon CloudWatch Agent to capture system logs and record them via Amazon CloudWatch Logs.

Answer: D

NEW QUESTION 27
A company is deploying a new application that uses Amazon EC2 instances. The company needs a solution to query application logs and AWS account API

activity. Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 31
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 35
A DevOps engineer is assisting with a multi-Region disaster recovery solution for a new application. The application consists of Amazon EC2 instances running in an Auto Scaling group and an Amazon Aurora MySQL DB cluster. The application must be available with an RTO of 120 minutes and an RPO of 60 minutes.
What is the MOST cost-effective way to meet these requirements?

A. Launch an Aurora DB cluster as an Aurora Replica in a different Regio
B. Create an AWS CloudFormation template for all compute resources and create a stack in two Region
C. Write a script that promotes the Aurora Replica to the primary instance in the event of a failure.
D. Launch an Aurora DB cluster as an Aurora Replica in a different Region and configure automatic cross-Region failove
E. Create an AWS CloudFormation template that includes an Auto Scaling group, and create a stack in two Region
F. Write a script that updates the CloudFormation stack in the disaster recovery Region to increase the number of instances.
G. Use AWS Lambda to create and copy a snapshot of the Aurora DB cluster to the destination Region hourl
H. Create an AWS CloudFormation template that includes an Auto Scaling group, and create a stack in two Region
I. Restore the Aurora DB cluster from a snapshot and update the Auto Scaling group to start launching instances.
J. Configure Amazon DynamoDB cross-Region replicatio
K. Create an AWS CloudFormation template that includes an Auto Scaling group, and create a stack in two Region
L. Write a script that will update the CloudFormation stack in the disaster recovery Region and promote the DynamoDB replica to the primary instance in the event of a failure.

Answer: D

NEW QUESTION 36
A company needs to introduce automatic DNS failover for a distributed web application to a disaster recovery or standby installation. The DevOps Engineer plans to configure Amazon Route 53 to provide DNS routing to alternate endpoint in the event of an application failure. What steps should the Engineer take to accomplish this? (Select TWO.)

A. Create Amazon Route 53 health checks for each endpoint that cannot be entered as alias record
B. Ensure firewall and routing rules allow Amazon Route 53 to send requests to the endpoints that are specified in the health checks.
C. Create alias records that route traffic to AWS resources and set the value of the Evaluate Target Health option to Yes, then create all the non-alias records.
D. Create a governing Amazon Route 53 record set, set it to failover, and associate it with the primary and secondary Amazon Route 53 record sets to distribute traffic to healthy DNS entries.
E. Create an Amazon CloudWatch alarm to monitor the primary Amazon Route 53 DNS entr
F. Then createan associated AWS Lambda function to execute the failover API call to Route 53 to the secondary DNS entry.

Answer: AC

NEW QUESTION 41
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 45
A company has built a web service that runs on Amazon EC2 instances behind an Application Load Balancer (ALB) the company has deployed the application in us-east-1 Amazon Route 53 provides an external DNS that routes traffic from example.com to the application, created with appropriate health checks.
The company has deployed a second environment for the application in eu-west-1 the company wants traffic to be routed to whichever environment results m the best response time for each user. If there is an outage in one Region, traffic should be directed to the other environment.
Which configuration will achieve this requirements?

A. •A subdomain us example com with weighted routing the US ALB with weight 2 and the EU ALB with weight 1•Another subdomain eu.example.com with weighted routing the EU ALB with weight 2 and the US ALU with weight 1•Geolocation routing records for example.com North America aliased to us example.com and Europe aliased to eu.example.com
B. •A subdomain us example com with latency-based routing the US ALB as the first target and the EU ALB as the second target.•Another subdomain eu.example.com with latency-based routin
C. The EU ALB as the first target and the US ALB as the second target.•Failover routing records for example.com aliased to us.example.com as the first target and eu.example.com as the second target.
D. •A subdomain us.example.com with failover routing the US ALB as primary and the EU ALB as secondary•Another subdomain eu.example.com with failover routing the EU ALB as primary and the US ALB as secondary•Latency-based routing records for example com that are aliased to us example com and eu.example.com
E. •A subdomain us.example.com with multivalue answer routin
F. the US ALB as first and the EU ALB as second•Another subdomain eu.example.com with failover routing the EU ALB as first and the US ALB as second•Failover routing records for example.com that are aliased to us.example.com and eu.example.com

Answer: B

NEW QUESTION 49
A DevOps engineer notices that all Amazon EC2 instances running behind an Application Load Balancer in an Auto Scaling group are failing to respond to user requests. The EC2 instances are also failing target group HTTP health checks.
Upon inspection, the engineer notices the application process was not running in any EC2 instances. There are a significant number of out of memory messages in the system logs. The engineer needs to improve the resilience of the application to cope with a potential application memory leak. Monitoring and notifications should be enabled to alert when there is an issue.
Which combination of actions will meet these requirements? {Select TWO.)

A. Change the Auto Scaling configuration to replace the instances when they fail the load balancer's health checks.
B. Change the target group health check HealthCheckIntervalSeconds parameter to reduce the interval between health checks.
C. Change the target group health checks from HTTP to TCP to check if the port where the application is listening is reachable.
D. Enable the available memory consumption metric within the Amazon CloudWatch dashboard for the entire Auto Scaling grou
E. Create an alarm when the memory utilization is hig
F. Associate an
G. Amazon SNS topic to the alarm to receive notifications when the alarm goes off.
H. Use the Amazon CloudWatch agent to collect the memory utilization of the EC2 instances in the Auto Scaling grou
I. Create an alarm when the memory utilization is high and associate an Amazon SNS topic to receive a notification.

Answer: BE

NEW QUESTION 52
A healthcare company has a critical application running in AWS. Recently, the company experienced some down time. if it happens again, the company needs to be able to recover its application in another AWS Region. The application uses Elastic Load Balancing and Amazon EC2 instances. The company also maintains a custom AMI that contains its application. This AMI is changed frequently.
The workload is required to run in the primary region, unless there is a regional service disruption, in which case traffic should fail over to the new region. Additionally, the cost for the second region needs to be low. The RTO is 2 hours.
Which solution allows the company to fail over to another region in the event of a failure, and also meet the above requirements?

A. Maintain a copy of the AMI from the main region in the backup regio
B. Create an Auto Scaling group with one instance using a launch configuration that contains the copied AM
C. Use an Amazon Route 53 record to direct traffic to the load balancer in the backup region in the event of failure, as require
D. Allow the Auto Scaling group to scale out as needed during a failure.
E. Automate the copying of the AMI in the main region to the backup regio
F. Generate an AWS Lambda function that will create an EC2 instance from the AMI and place it behind a load balance
G. Using the same Lambda function, point the Amazon Route 53 record to the load balancer in the backup regio
H. Trigger the Lambda function in the event of a failure.
I. Place the AMI in a replicated Amazon S3 bucke
J. Generate an AWS Lambda function that can create a launch configuration and assign it to an already created Auto Scaling grou
K. Have one instance in this Auto Scaling group ready to accept traffi
L. Trigger the Lambda function in the event of a failur
M. Use an Amazon Route 53 record and modify it with the same Lambda function to point to the load balancer in the backup region.
N. Automate the copying of the AMI to the backup regio
O. Create an AWS Lambda function that can create a launch configuration and assign it to an already created Auto Scaling grou
P. Set the Auto Scaling group maximum size to 0 and only increase it with the Lambda function during a failur
Q. Trigger the Lambda function in the event of a failur
R. Use an Amazon Route 53 record and modify it with the same Lambda function to point to the load balancer in the backup region.

Answer: C

NEW QUESTION 55

A DevOps engineer is researching the least expensive way to implement an image batch processing cluster on AWS. The application cannot run in Docker containers and must run on Amazon EC2. The batch job stores checkpoint data on an NFS and can tolerate interruptions. Configuring the cluster software from a generic EC2 Linux image takes 30 minutes.
What is the MOST cost-effective solution?

A. Use Amazon EFS for checkpoint dat
B. To complete the jo
C. use an EC2 Auto Scaling group and an On-Demand pricing model to provision EC2 instances temporarily.
D. Use GlusterFS on EC2 instances for checkpoint dat
E. To run the batch jo
F. configure EC2 instances manuall
G. When the job completes, shut down the instances manually.
H. Use Amazon EFS for checkpoint dat
I. Use EC2 Fleet to launch EC2 Spot Instances, and utilize user data to configure the EC2 Linux instance on startup.
J. Use Amazon EFS for checkpoint dat
K. Use EC2 Fleet to launch EC2 Spot Instance
L. Create a custom AMI for the cluster and use the latest AMI when creating instances.

Answer: A

NEW QUESTION 57
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 59
A company is using AWS CodeDeploy to manage its application deployments. Recently, the Development team decided to use GitHub for version control, and the team is looking for ways to integrate the GitHub repository with CodeDeploy. The team also needs to develop a way to automate deployment whenever there is a new commit on that repository. The team is currently deploying new application revisions by manually indicating the Amazon S3 location.
How can the integration be achieved in the MOST efficient way?

A. Create a GitHub webhook to replicate the repository to AWS CodeCommi
B. Create an AWSCodePipeline pipeline that uses CodeCommit as a source provider and AWS CodeDeploy as a deployment provide
C. Once configured, commit a change to the GitHub repository to start the first deployment.
D. Create an AWS CodePipeline pipeline that uses GitHub as a source provider and AWS CodeDeploy as a deployment provide
E. Connect this new pipeline with the GitHub account and instruct CodePipeline to use webhooks in GitHub to automatically start the pipeline when a change occurs.
F. Create an AWS Lambda function to check periodically if there has been a new commit within the GitHub repositor
G. If a new commit is found, trigger a CreateDeployment API call to AWS CodeDeploy to start a new deployment based on the last commit ID within the deployment group.
H. Create an AWS CodeDeploy custom deployment configuration to associate the GitHub repository with the deployment grou
I. During the association process, authenticate the deployment group with GitHub to obtain the GitHub security authentication toke
J. Configure the deployment group options to automatically deploy if a new commit is foun
K. Perform a new commit to the GitHub repository to trigger the first deployment.

Answer: B

NEW QUESTION 62
A company discovers that some IAM users have been storing their AWS access keys in configuration files that have been pushed to a Git repository hosting service.
Which solution will require the LEAST amount of management overhead while preventing the exposed AWS access keys from being used?

A. Build an application that will create a list of all AWS access keys in the account and search each key on Git repository hosting service
B. If a match is found, configure the application to disable the associated access ke
C. Then deploy the application to an AWS Elastic Beanstalk worker environment and define a periodic task to invoke the application every hour.
D. Use Amazon Inspector to detect when a key has been exposed onlin
E. Have Amazon Inspector send a notification to an Amazon SNS topic when a key has been expose
F. Create an AWS Lambda function subscribed to the SNS topic to disable the IAM user to whom the key belongs, and then delete the key so that it cannot be used.
G. Configure AWS Trusted Advisor and create an Amazon CloudWatch Events rule that uses Trusted Advisor as the event sourc
H. Configure the CloudWatch Events rule to invoke an AWS Lambda function as the targe
I. If the Lambda function finds the exposed access keys, then have it disable the access key so that it cannot be used.
J. Create an AWS Config rule to detect when a key is exposed onlin
K. Haw AWS Config send change notifications to an SNS topi
L. Configure an AWS Lambda function that is subscribed to the SNS topic to check the notification sent by AWS Config, and then disable the access key so it cannot be used.

Answer: C

Explanation:

https://github.com/aws/Trusted-Advisor-Tools/tree/master/ExposedAccessKeys/stepbystep

NEW QUESTION 64
A company runs a production application workload in a single AWS account that uses Amazon Route 53, AWS Elastic Beanstalk, and Amazon RDS. In the event of a security incident, the Security team wants the application workload to fail over to a new AWS account. The Security team also wants to block all access to the original account immediately, with no access to any AWS resources in the original AWS account, during forensic analysis.
What is the most cost-effective way to prepare to fail over to the second account prior to a security incident?

A. Migrate the Amazon Route 53 configuration to a dedicated AWS accoun
B. Mirror the Elastic Beanstalk configuration in a different accoun
C. Enable RDS Database Read Replicas in a different account.
D. Migrate the Amazon Route 53 configuration to a dedicated AWS accoun
E. Save/copy the Elastic Beanstalk configuration files in a different AWS accoun
F. Copy snapshots of the RDS Database to a different account.
G. Save/copy the Amazon Route 53 configurations for use in a different AWS account after an incident.Save/copy Elastic Beanstalk configuration files to a different accoun
H. Enable the RDS database read replica in a different account.
I. Save/copy the Amazon Route 53 configurations for use in a different AWS account after an incident.Mirror the configuration of Elastic Beanstalk in a different accoun
J. Copy snapshots of the RDS database to a different account.

Answer: B

NEW QUESTION 68
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 70
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 74
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:

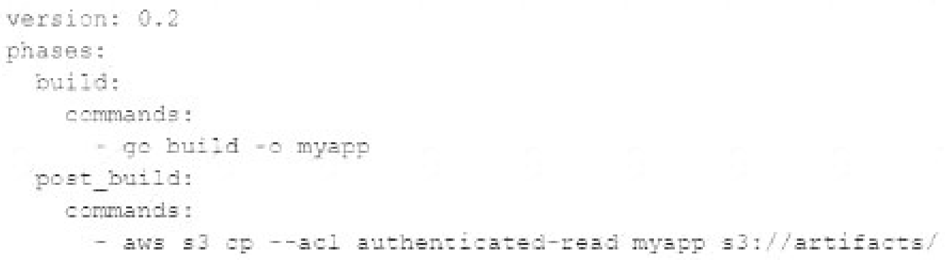

```
version: 0.2
phases:
  build:
    commands:
      - go build -o myapp
  post_build:
    commands:
      - aws s3 cp --acl authenticated-read myapp s3://artifacts/
```

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 75
A company runs an application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones in us-east1. The application stores data in an Amazon RDS MySQL Multi-AZ DB instance.
A DevOps Engineer wants to modify the current solution and create a hot standby of the environment in another region to minimize downtime if a problem occurs in us-east-1.
Which combination of steps should the DevOps Engineer take to meet these requirements? (Select THREE.)

A. Add a health check to the Amazon Route 53 alias record to evaluate the health of the primary region.Use AWS Lambda, configured with an Amazon CloudWatch Events trigger, to elect the Amazon RDS master in the disaster recovery region.
B. Create a new Application Load Balancer and Auto Scaling group in the disaster recovery region.
C. Extend the current Auto Scaling group to the subnets in the disaster recovery region.
D. Enable multi-region failover for the RDS configuration for the database instance.
E. Deploy a read replica of the RDS instance in the disaster recovery region.
F. Create an AWS Lambda function to evaluate the health of the primary regio
G. If it fails, modify the Amazon Route 53 record to point at the disaster recovery region and elect the RDS master.

Answer: ABE

Explanation:
https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazon-rds/

NEW QUESTION 78
A DevOps team needs to query information in application logs that are generated by an application running multiple Amazon EC2 instances deployed with AWS Elastic Beanstalk.
Instance log streaming to Amazon CloudWatch Logs was enabled on Elastic Beanstalk. Which approach would be the MOST cost-efficient?

A. Use a CloudWatch Logs subscription to trigger an AWS Lambda function to send the log data to an Amazon Kinesis Data Firehouse stream that has an Amazon S3 bucket destinatio
B. Use Amazon Athena to query the log data from the bucket.
C. Use a CloudWatch Logs subscription to trigger an AWS Lambda function to send the log data to an Amazon Kinesis Data Firehouse stream that has an Amazon S3 bucket destinatio
D. Use a new Amazon Redshift cluster and Amazon Redshift Spectrum to query the log data from the bucket.
E. Use a CloudWatch Logs subscription to send the log data to an Amazon Kinesis Data Firehouse stream that has an Amazon S3 bucket destinatio
F. Use Amazon Athena to query the log data from the bucket.
G. Use a CloudWatch Logs subscription to send the log data to an Amazon Kinesis Data Firehouse stream that has an Amazon S3 bucket destinatio
H. Use a new Amazon Redshift cluster and Amazon Redshift Spectrum to query the log data from the bucket.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/Subscriptions.html

NEW QUESTION 83
After a recent audit, a company decided to implement a new disaster recovery strategy for Its Amazon S3 data and its MySQL database running on Amazon EC2.
Management wants the ability to recover to a secondary AWS Region with an RPO under 5 seconds and a RTO under 1 minute.
Which actions will meet the requirements while MINIMIZING operational overhead? (Select TWO.)

A. Modify the application to write to both Regions at the same time when uploading objects to Amazon S3
B. Migrate the database to an Amazon Aurora multi-master in the primary and secondary Regions.
C. Migrate the database to Amazon RDS with a read replica in the secondary Region
D. Migrate to Amazon Aurora Global Database.
E. Set up S3 cross-Region replication with a replication SLA for the S3 buckets where objects are being put.

Answer: AE

NEW QUESTION 87
A company has multiple development teams sharing one AWS account. The development team's manager wants to be able to automatically stop Amazon EC2 instances and receive notifications if resources are idle and not tagged as production resources
Which solution will meet these requirements?

A. Use a scheduled Amazon CloudWatch Events rule to filter for Amazon EC2 instance status checks and identify idle EC2 instance

B. Use the CloudWatch Events rule to target an AWS Lambda function to stop non-production instances and send notifications.
C. Use a scheduled Amazon CloudWatch Events rule to filter AWS Systems Manager events and identifyidle EC2 instances and resource
D. Use the CloudWatch Events rule to target an AWS Lambda function to stop non-production instances and send notifications.
E. Use a scheduled Amazon CloudWatch Events rule to target a custom AWS Lambda function that runs AWS Trusted Advisor checks Create a second CloudWatch Events rule to filter events from Trusted Advisor to trigger a Lambda function to stop idle non-production instances and send notifications
F. Use a scheduled Amazon CloudWatch Events rule to target Amazon Inspector events for idle EC2 instances Use the CloudWatch Events rule to target the AWS Lambda function to stop non-production instances and send notifications

Answer: A

NEW QUESTION 89
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 91
An application's users ate encountering bugs immediately after Amazon API Gateway deployments. The development team deploys once or twice a day and uses a blue/green deployment strategy with custom health checks and automated rollbacks. The team wants to limit the number of users affected by deployment bugs and receive notifications when rollbacks are needed.
Which combination of steps should a DevOps engineer use to meet these requests? (Select TWO.)

A. Implement a blue/green strategy using path mappings.
B. Implement a canary deployment strategy.
C. Implement a rolling deployment strategy using multiple stages.
D. Use Amazon CloudWatch alarms to notify the development team.
E. Use Amazon CloudWatch Events to notify the development team.

Answer: BD

NEW QUESTION 96
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 98
A Development team is adding a new country to an e-commerce application. This addition requires that new application features be added to the shipping component of the application. The team has not decided if all new features should be added, as some will take approximately six weeks to build. While the final decision on the shipping component features is being made, other team members are continuing to work on other features of the application.
Based on this situation, how should the application feature deployments be managed?

A. Add the code updates as commits to the release branc
B. The team can delay the deployment until all features are ready.
C. Add the code updates as commits to a feature branc
D. Merge the commits to a release branch as features are ready.
E. Add the code updates as a single commit when a feature is read
F. Tag this commit with "new-country."

G. Create a new repository named "new-country". Commit all the code changes to the new repository.

Answer: B

NEW QUESTION 101
A company's application is running on Amazon EC2 instances in an Auto Scaling group. A DevOps engineer needs to ensure there are at least four application servers running at all times. Whenever an update has to be made to the application, the engineer creates a new AMI with the updated configuration and updates the AWS CloudFormation template with the new AMI ID. After the stack update finishes, the engineer manually terminates the old instances one by one. verifying that the new instance is operational before proceeding. The engineer needs to automate this process.
Which action will allow for the LEAST number of manual steps moving forward?

A. Update the CloudFormation template to include the UpdatePolicy attribute with the AutoScalingRollingUpdate policy.
B. Update the CloudFormation template to include the UpdatePolicy attribute with the AutoScalingReplacingUpdate policy.
C. Use an Auto Scaling lifecycle hook to verify that the previous instance is operational before allowing the DevOps engineer's selected instance to terminate.
D. Use an Auto Scaling lifecycle hook to confirm there are at least four running instances before allowing the DevOps engineer's selected instance to terminate.

Answer: A

NEW QUESTION 105
An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.
When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.
How should the company meet these requirements with the LEAST amount of application changes?

A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.
B. Use Amazon DynamoDB global tables for the product catalog and regional tables for the customer information and purchases
C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.
D. Use Aurora for the product catalog and Amazon DynamoDB global tables for the customer information and purchases.

Answer: C

NEW QUESTION 106
An application is deployed on Amazon EC2 instances running in an Auto Scaling group. During the bootstrapping process, the instances register their private IP addresses with a monitoring system. The monitoring system performs health checks frequently by sending ping requests to those IP addresses and sending alerts if an instance becomes non-responsive.
The existing deployment strategy replaces the current EC2 instances with new ones. A DevOps engineer has noticed that the monitoring system is sending false alarms during a deployment, and is tasked with stopping these false alarms.
Which solution will meet these requirements without affecting the current deployment method?

A. Define an Amazon CloudWatch Events target, an AWS Lambda function, and a lifecycle hook attached to the Auto Scaling grou
B. Configure CloudWatch Events to invoke Amazon SNS to send a message to the systems administrator group for remediation.
C. Define an AWS Lambda function and a lifecycle hook attached to the Auto Scaling grou
D. Configure the lifecycle hook to invoke the Lambda function, which removes the entry of the private IP from the monitoring system upon instance termination.
E. Define an Amazon CloudWatch Events target, an AWS Lambda function, and a lifecycle hook attached to the Auto Scaling grou
F. Configure CloudWatch Events to invoke the Lambda function, which removes the entry of the private IP from the monitoring system upon instance termination.
G. Define an AWS Lambda function that will run a script when instance termination occurs in an Auto Scaling grou
H. The script will remove the entry of the private IP from the monitoring system.

Answer: C

NEW QUESTION 110
A company has thousands of Amazon EC2 instances as well as hundreds of virtual machines on-premises. Developers routinely sign in to the console for on-premises systems to perform troubleshooting. The developers want to sign in to AWS instances to run performance tools, but are unable to due to the lack of a central console logging system. A DevOps engineer wants to ensure that console access is logged on all systems.
Which combination of steps will meet these requirements? (Select TWO.)

A. Attach a role to all AWS instances that contains the appropriate permission
B. Create an AWS Systems Manager managed-instance activatio
C. Install and configure Systems Manager Agent on on-premises machines.
D. Enable AWS Systems Manager Session Manager logging to an Amazon S3 bucke
E. Direct developers to connect to the systems with Session Manager only.
F. Enable AWS Systems Manager Session Manager logging to AWS CloudTrai
G. Direct developers to continue normal sign-in procedures for on-premise
H. Use Session Manager for AWS instances.
I. Install and configure an Amazon CloudWatch Logs agent on all system
J. Create an AWS Systems Manager managed-instance activation.
K. Set up a Site-to-Site VPN connection between the on-premises and AWS network
L. Set up a bastion instance to allow developers to sign in to the AWS instances.

Answer: AB

NEW QUESTION 112
An IT department manages a portfolio with Windows and Linux (Amazon and Red Hat Enterprise Linux) servers both on-premises and on AWS. An audit reveals that there is no process for updating OS and core application patches, and that the servers have inconsistent patch levels.
Which of the following provides the MOST reliable and consistent mechanism for updating and maintaining all servers at the recent OS and core application patch levels?

A. Install AWS Systems Manager agent on all on-premises and AWS server
B. Create Systems Manager Resource Group
C. Use Systems Manager Patch Manager with a preconfigured patch baseline to run scheduled patch updates during maintenance windows.
D. Install the AWS OpsWorks agent on all on-premises and AWS server
E. Create an OpsWorks stack with separate layers for each operating system, and get a recipe from the Chef supermarket to run the patch commands for each layer during maintenance windows.
F. Use a shell script to install the latest OS patches on the Linux servers using yum and schedule it to run automatically using cro
G. Use Windows Update to automatically patch Windows servers.
H. Use AWS Systems Manager Parameter Store to securely store credentials for each Linux and Windows serve
I. Create Systems Manager Resource Group
J. Use the Systems Manager Run Command to remotely deploy patch updates using the credentials in Systems Manager Parameter Store

Answer: A

Explanation:
1- https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-patch-patchgroups.html 2- https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html

NEW QUESTION 116
A DevOps engineer has automated a web service deployment using AWS CodePipelme with the following steps:
• An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
• An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
• A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment The quality assurance (QA) team has asked for permission to inspect the build artifact before the deployment to the production environment occurs. The OA team wants to run an internal automated penetration testing tool (invoked using a REST API call) to run some manual tests.
Which combination of actions will fulfill this request? (Select TWO.)

A. Insert a manual approval action between the test and deployment actions of Jtue pipeline.
B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.
C. Update the CodeDeploy deployment group so it requires manual approval to proceed
D. Update the pipeline to directly trigger the REST API for the automated penetration testing tool.
E. Update the pipeline to invoke a Lambda function that triggers the REST API for the automated penetration testing tool.

Answer: BD

NEW QUESTION 118
A company gives its employees limited rights to AWS. DevOps engineers have the ability to assume an administrator role. For tracking purposes, the security team wants to receive a near-real-time notification when the administrator role is assumed.
How should this be accomplished?

A. Configure AWS Config to publish logs to an Amazon S3 bucke
B. Use Amazon Athena to query the logs and send a notification to the security team when the administrator role is assumed.
C. Configure Amazon GuardDuty to monitor when the administrator role is assumed and send a notification to the security team.
D. Create an Amazon EventBridge (Amazon CloudWatch Events) event rule using an AWS Management Console sign-in events event pattern that publishes a message to an Amazon SNS topic if the administrator role is assume
E. [^
F. Create an Amazon EventBridge (Amazon CloudWatch Events) events rule using an AWS API call thatuses an AWS CloudTrail event pattern to trigger an AWS Lambda function that publishes a message to an Amazon SNS topic if the administrator role is assumed.

Answer: D

NEW QUESTION 119
A company wants to adopt a methodology for handling security threats from leaked and compromised IAM access keys. The DevOps Engineer has been asked to automate the process of acting upon compromised access keys, which includes identifying users, revoking their permissions, and sending a notification to the Security team.
Which of the following would achieve this goal?

A. Use the AWS Trusted Advisor generated security report for access key
B. Use Amazon EMR to run analytics on the repor
C. Identify compromised IAM access keys and delete the
D. Use Amazon CloudWatch with an EMR Cluster State Change event to notify the Security team.
E. Use AWS Trusted Advisor to identify compromised access key
F. Create an Amazon CloudWatch Events rule with Trusted Advisor as the event source, and AWS Lambda and Amazon SNS as target
G. Use AWS Lambda to delete compromised IAM access keys and Amazon SNS to notify the Security team.
H. Use the AWS Trusted Advisor generated security report for access key
I. Use AWS Lambda to scan through the repor
J. Use scan result inside AWS Lambda and delete compromised IAM access key
K. Use Amazon SNS to notify the Security team.
L. Use AWS Lambda with a third-party library to scan for compromised access key
M. Use scan result inside AWS Lambda and delete compromised IAM access key
N. Create Amazon CloudWatch custom metrics for compromised key
O. Create a CloudWatch alarm on the metrics to notify the Security team.

Answer: B

NEW QUESTION 120
A company wants to use a grid system for a proprietary enterprise in-memory data store on top of AWS. This system can run in multiple server nodes in any Linux-based distribution. The system must be able to reconfigure the entire cluster every time a node is added or removed. When adding or removing nodes, an / etc./cluster/nodes.config file must be updated, listing the IP addresses of the current node members of that cluster

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

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

Answer: A

Explanation:
https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html

NEW QUESTION 121
A company is adopting serverless computing and is migrating some of its existing applications to AWS Lambda A DevOps engineer must come up with an automated deployment strategy using AWS CodePipeline that should include proper version controls, branching strategies, and rollback methods
Which combination of steps should the DevOps engineer follow when setting up the pipeline? (Select THREE)

A. Use Amazon S3 as the source code repository
B. Use AWS CodeCommit as the source code repository
C. Use AWS CloudFormation to create an AWS Serverless Application Model (AWS SAM) template for deployment.
D. Use AWS CodeBuild to create an AWS Serverless Application Model (AWS SAM) template for deployment
E. Use AWS CloudFormation to deploy the application
F. Use AWS CodeDeploy to deploy the application.

Answer: ABC

NEW QUESTION 126
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 128
A DevOps Engineer is deploying a new web application. The company chooses AWS Elastic Beanstalk for deploying and managing the web application, and Amazon RDS MySQL to handle persistent data. The company requires that new deployments have minimal impact if they fail. The application resources must be at full capacity during deployment, and rolling back a deployment must also be possible.
Which deployment sequence will meet these requirements?

A. Deploy the application using Elastic Beanstalk and connect to an external RDS MySQL instance using Elastic Beanstalk environment propertie
B. Use Elastic Beanstalk features for a blue/green deployment to deploy the new release to a separate environment, and then swap the CNAME in the two environments to redirect traffic to the new version.
C. Deploy the application using Elastic Beanstalk, and include RDS MySQL as part of the environment.Use default Elastic Beanstalk behavior to deploy changes to the application, and let rolling updates deploy changes to the application.
D. Deploy the application using Elastic Beanstalk, and include RDS MySQL as part of the environment.Use Elastic Beanstalk immutable updates for application deployments.
E. Deploy the application using Elastic Beanstalk, and connect to an external RDS MySQL instance using Elastic Beanstalk environment propertie
F. Use Elastic Beanstalk immutable updates for application deployments.

Answer: A

NEW QUESTION 131
A DevOps Engineer manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The Engineer needs to implement a deployment strategy that:
Launches a second fleet of instances with the same capacity as the original fleet. Maintains the original fleet unchanged while the second fleet is launched. Transitions traffic to the second fleet when the second fleet is fully deployed. Terminates the original fleet automatically 1 hour after transition.
Which solution will satisfy these requirements?

A. Use an AWS CloudFormation template with a retention policy for the ALB set to 1 hou
B. Update the Amazon Route 53 record to reflect the new ALB.
C. Use two AWS Elastic Beanstalk environments to perform a blue/green deployment from the original environment to the new on
D. Create an application version lifecycle policy to terminate the original environment in 1 hour.
E. Use AWS CodeDeploy with a deployment group configured with a blue/green deployment configuratio
F. Select the option Terminate the original instances in the deployment group with a waiting period of 1 hour.
G. Use AWS Elastic Beanstalk with the configuration set to Immutabl
H. Create an .ebextension using the Resources key that sets the deletion policy of the ALB to 1 hour, and deploy the application.

Answer: B

NEW QUESTION 133
A web application for healthcare services runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. A DevOps Engineer must create a mechanism in which an EC2 instance can be taken out of production so its system logs can be analyzed for issues to quickly troubleshot problems on the web tier.
How can the Engineer accomplish this task while ensuring availability and minimizing downtime?

A. Implement EC2 Auto Scaling groups cooldown period
B. Use EC2 instance metadata to determine the instance state, and an AWS Lambda function to snapshot Amazon EBS volumes to preserve system logs.
C. Implement Amazon CloudWatch Events rule
D. Create an AWS Lambda function that can react to an instance termination to deploy the CloudWatch Logs agent to upload the system and access logs to Amazon S3 for analysis.
E. Terminate the EC2 instances manuall
F. The Auto Scaling service will upload all log information toCloudWatch Logs for analysis prior to instance termination.
G. Implement EC2 Auto Scaling groups with lifecycle hook
H. Create an AWS Lambda function that can modify an EC2 instance lifecycle hook into a standby state, extract logs from the instance through a remote script execution, and place them in an Amazon S3 bucket for analysis.

Answer: D

NEW QUESTION 136
During metric analysis, your team has determined that the company's website during peak hours is experiencing response times higher than anticipated. You currently rely on Auto Scaling to make sure that you are scaling your environment during peak windows. How can you improve your Auto Scaling policy to reduce this high response time? Choose 2 answers.

A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto Scaling policy to have betterfine-grain insight.
B. IncreaseyourAutoScalinggroup'snumberofmaxservers.
C. Create a script that runs and monitors your servers; when it detects an anomaly in load, it posts to an Amazon SNS topic that triggers Elastic Load Balancing to add more servers to the load balancer.
D. Push custom metrics to CloudWatch for your application that include more detailed information about your web application, such as how many requests it is handling and how many are waiting to be processed.

Answer: BD

Explanation:
Option B makes sense because maybe the max servers is low hence the application cannot handle the peak load.
Option D helps in ensuring Autoscaling can scale the group on the right metrics.
For more information on Autoscaling health checks, please refer to the below document link: from AWS
- http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html

NEW QUESTION 140
A company uses AWS CodePipeline to manage and deploy infrastructure as code. The infrastructure is defined in AWS CloudFormation templates and is primarily comprised of multiple Amazon EC2 instances and Amazon RDS databases. The Security team has observed many operators creating inbound security group rules with a source CIDR of 0 0 0 0/0 and would like to proactively stop the deployment of rules with open CIDRs
The DevOps Engineer will implement a predeptoyment step that runs some security checks over the CloudFormation template before the pipeline processes it. This check should allow only inbound security group rules with a source CIDR of 0.0.0.0/0 if the rule has the description "Security Approval Ref XXXXX (where XXXXX is a preallocated reference). The pipeline step should fail if this condition is not met and the deployment should be blocked
How should this be accomplished?

A. Enable a SCP in AWS Organization
B. The policy should deny access to the API call Create Security GroupRule if the rule specifies 0.0.0.0/0 without a description referencing a security approval
C. Add an initial stage to CodePipeline called Security Chec
D. This stage should call an AWS Lambda function that scans the CloudFormation template and fails the pipeline if it finds 0.0.0.0/0 in a security group without a description referencing a security approval
E. Create an AWS Config rule that is triggered on creation or edit of resource type EC2 SecurityGroup.This rule should call an AWS Lambda function to send a failure notification if the security group has any rules with a source CIDR of 0.0.0.0/0 without a description referencing a security approval.
F. Modify the IAM role used by CodePipelin
G. The IAM policy should deny access.

Answer: B

NEW QUESTION 141
A large enterprise is deploying a web application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS Oracle DB instance and Amazon DynamoDB. There are separate environments for development, testing, and production.
What is the MOST secure and flexible way to obtain password credentials during deployment?

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

Answer: B

Explanation:
https://www.1strategy.com/blog/2019/02/28/aws-parameter-store-vs-aws-secrets-manager/

NEW QUESTION 142
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 145
You have decided that you need to change the instance type of your production instances which are running as part of an AutoScaling group. The entire architecture is deployed using CloudFormation Template. You currently have 4 instances in Production. You cannot have any interruption in service and need to ensure 2 instances are always runningduring the update? Which of the options below listed can be used for this?

A. AutoScalingRollingUpdate
B. AutoScalingScheduledAction
C. AutoScalingReplacingUpdate
D. AutoScalingIntegrationUpdate

Answer: A

Explanation:
The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePoIicy attribute. This is used to define how an Auto Scalinggroup resource is updated when an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified. For more information on Autoscaling updates, please refer to the below link:
- https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/

NEW QUESTION 146
A DevOps engineer is currently running a container-based workload on-premises The engineer wants to move the application to AWS, but needs to keep the on-premises solution active because not all APIs will move at the same time. The traffic between AWS and the on-premises network should be secure and encrypted at all times. Low management overload is also a requirement.
Which combination of actions will meet these criteria? (Select THREE.)

A. Create a Network Load Balancer an
B. for each service, create a listener that points to the correct set of containers either in AWS or on-premises.
C. Create an Application Load Balancer and, for each service, create a listener that points to the correct set of containers either in AWS or on-premises.
D. Host the AWS containers in Amazon ECS with an EC2 launch type.
E. Host the AWS containers in Amazon ECS with a Fargate launch type
F. Use Amazon API Gateway to front the workload, and create a VPC link so API Gateway can forward API calls to the on-premises network through a VPN connection.
G. Use Amazon API Gateway to front the workload, and set up public endpoints for the on-premises APIs so API Gateway can access them.

Answer: BDF

NEW QUESTION 151
A DevOps engineer used an AWS CloudFormation custom resource to set up AD Connector. The AWS Lambda function executed and created AD Connector, but CloudFormation is not transitioning from CREATE_IN_PROGRESS to CREATE.COMPLETE.
Which action should the engineer take to resolve this issue?

A. Ensure the Lambda function code has exiled successfully.
B. Ensure the Lambda function code returns a response to the pre-signed URL.
C. Ensure the Lambda function IAM role has cloudformation:UpdateStack permissions for the stack ARN.
D. Ensure the Lambda function IAM role has ds:ConnectDirectory permissions for the AWS account.

Answer: A

NEW QUESTION 152
A DevOps Engineer at a startup cloud-based gaming company has the task formalizing deployment strategies. The strategies must meet the following requirements:
Use standard Git commands, such as git clone and git push for the code repository. Management tools should maximize the use of platform solutions where possible. Deployment packages must be immutable and in the form of Docker images.
How can the Engineer meet these requirements?

A. Use AWS CodePipeline to trigger a build process when software is pushed to a self-hosted GitHub repositor
B. CodePipeline will use a Jenkins build server to build new Docker image
C. CodePipeline will deploy into a second target group in Amazon ECS behind an Application Load Balance
D. Cutover will be managed by swapping the listener rules on the Application Load Balancer.
E. Use AWS CodePipeline to trigger a build process when software is pushed to a private GitHub repositor
F. CodePipeline will use AWS CodeBuild to build new Docker image
G. CodePipeline will deploy into a second target group in Amazon ECS behind an Application Load Balance
H. Cutover will be managed by swapping the listener rules on the Application Load Balancer.
I. Use a Jenkins pipeline to trigger a build process when software is pushed to a private GitHub repository.AWS CodePipeline will use AWS CodeBuild new Docker image
J. CodePipeline will deploy into a second target group in Amazon ECS behind an Application Load Balance
K. Cutover will be managed by swapping the listener rules on the Application Load Balancer.
L. Use AWS CodePipeline to trigger a build process when software is pushed to an AWS CodeCommit repository CodePipeline will use an AWS CodeBuild build server to build new Docker image
M. CodePipeline will deploy into a second target group in a Kubernetes Cluster hosted on Amazon EC2 behind an Application Load Balance
N. Cutover will be managed by swapping the listener rules on the Application Load Balancer.

Answer: B

NEW QUESTION 154
A company develops and maintains a web application using Amazon EC2 instances and an Amazon RDS for SQL Server DB instance in a single Availability Zone The resources need to run only when new deployments are being tested using AWS CodePipeline. Testing occurs one or more times a week and each test takes 2-3 hours to run. A DovOps engineer wants a solution that does not change the architecture components.
Which solution will meet these requirements in the MOST cost-effective manner?

A. Convert the RDS database to an Amazon Aurora Serverless database Use an AWS Lambda function to start and stop the EC2 instances before and after tests
B. Put the EC2 instances into an Auto Scaling grou
C. Schedule scaling to run at the start of the deployment tests.
D. Replace the EC2 instances with EC2 Spot Instances and the RDS database with an RDS Reserved Instance.
E. Subscribe Amazon CloudWatch Events to CodePipeline to trigger AWS Systems Manager Automation documents that start and stop all EC2 and RDS instances before and after deployment tests.

Answer: A

NEW QUESTION 156
A DevOps Engineer is working on a project that is hosted on Amazon Linux and has failed a security review. The DevOps Manager has been asked to review the company buildspec.yaml file for an AWS CodeBuild project and provide recommendations. The buildspec.yaml file is configured as follows:

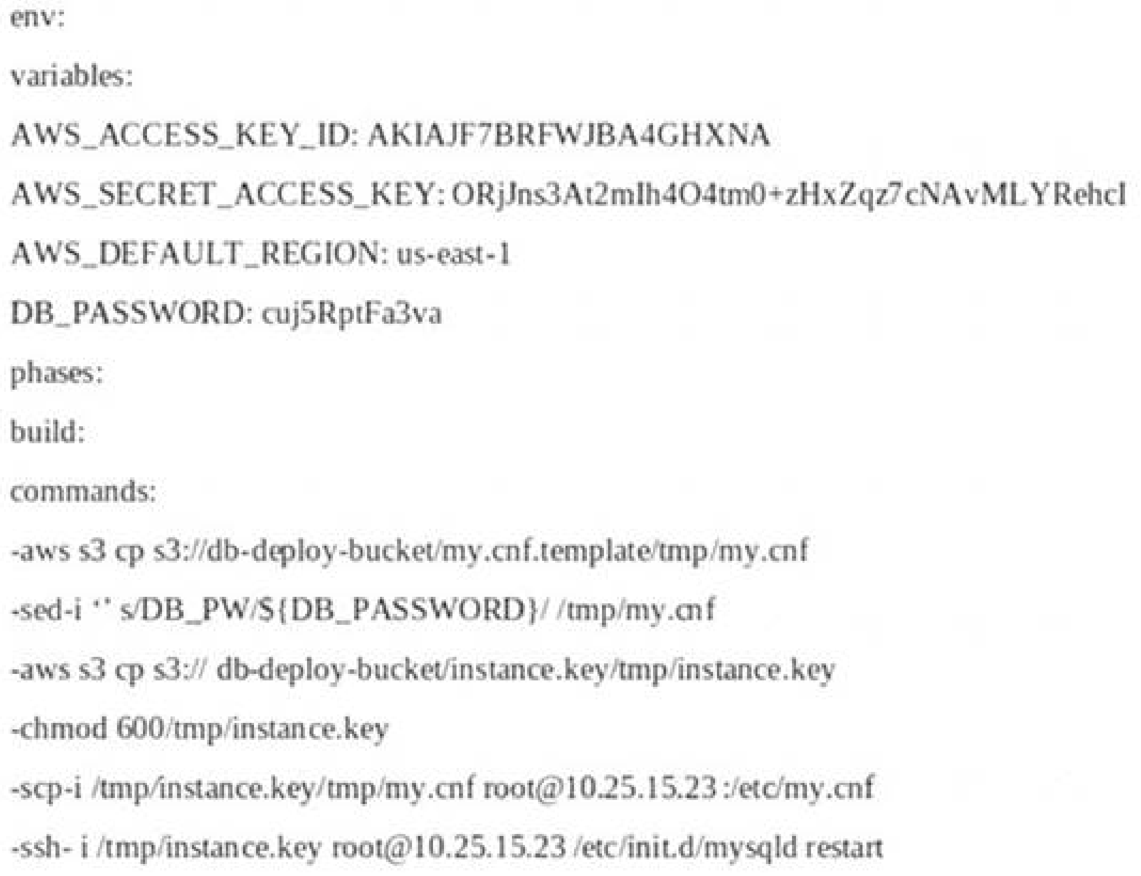

```
env:
variables:
AWS_ACCESS_KEY_ID: AKIAJF7BRFWJBA4GHXNA
AWS_SECRET_ACCESS_KEY: ORjJns3At2mlh4O4tm0+zHxZqz7cNAvMLYRehcl
AWS_DEFAULT_REGION: us-east-1
DB_PASSWORD: cuj5RptFa3va
phases:
build:
commands:
-aws s3 cp s3://db-deploy-bucket/my.cnf.template/tmp/my.cnf
-sed-i '' s/DB_PW/${DB_PASSWORD}/ /tmp/my.cnf
-aws s3 cp s3:// db-deploy-bucket/instance.key/tmp/instance.key
-chmod 600/tmp/instance.key
-scp-i /tmp/instance.key/tmp/my.cnf root@10.25.15.23 :/etc/my.cnf
-ssh- i /tmp/instance.key root@10.25.15.23 /etc/init.d/mysqld restart
```

What changes should be recommended to comply with AWS security best practices? (Select THREE.)

A. Add a post-build command to remove the temporary files from the container before termination to ensure they cannot be seen by other CodeBuild users.
B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable.
C. Store the DB_PASSWORD as a SecureString value in AWS Systems Manager Parameter Store and then remove the DB_PASSWORD from the environment variables.
D. Move the environment variables to the "˜db-deploy-bucket' Amazon S3 bucket, add a prebuild stage to download, then export the variables.
E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.

F. Scramble the environment variables using XOR followed by Base64, add a section to install, and then run XOR and Base64 to the build phase.

Answer: BCE

Explanation:
https://aws.amazon.com/codebuild/faqs/

NEW QUESTION 158
A company that runs many workloads on AWS has an Amazon EBS spend that has increased over time. The DevOps team notices there are many unattached EBS volumes. Although there are workloads where volumes are detached, volumes over 14 days old are stale and no longer needed. A DevOps engineer has been tasked with creating automation that deletes unattached EBS volumes that have been unattached for 14 days.
Which solution will accomplish this?

A. Configure the AWS Config ec2-vo!ume-inuse-check managed rule with a configuration changes trigger type and an Amazon EC2 volume resource targe
B. Create a new Amazon CloudWatch Events rule scheduled to execute an AWS Lambda function in 14 days to delete the specified EBS volume.
C. Use Amazon EC2 and Amazon Data Lifecycle Manager to configure a volume lifecycle polic
D. Set the interval period for unattached EBS volumes to 14 days and set the retention rule to delet
E. Set the policy target volumes as
F. Create an Amazon CloudWatch Events rule to execute an AWS Lambda function dail
G. The Lambda function should find unattached EBS volumes and tag them with the current date, and delete unattached volumes that have tags with dates that are more than 14 days old.
H. Use AWS Trusted Advisor to detect EBS volumes that have been detached for more than 14 days.Execute an AWS Lambda function that creates a snapshot and then deletes the EBS volume.

Answer: C

NEW QUESTION 161
You have an ELB setup in AWS with EC2 instances running behind it. You have been requested to monitor the incoming connections to the ELB. Which of the below options can suffice this requirement?

A. UseAWSCloudTrail with your load balancer
B. Enable access logs on the load balancer
C. Use a CloudWatch Logs Agent
D. Create a custom metric CloudWatch filter on your load balancer

Answer: B

Explanation:
Clastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Cach log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues.
Option A is invalid because this service will monitor all AWS services Option C and D are invalid since CLB already provides a logging feature.
For more information on ELB access logs, please refer to the below document link: from AWS
> http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html

NEW QUESTION 162
An online company uses Amazon EC2 Auto Scaling extensively to provide an excellent customer experience while minimizing the number of running EC2 instances. The company's self-hosted Puppet environment in the application layer manages the configuration of the instances. The IT manager wants the lowest licensing costs and wants to ensure that whenever the EC2 Auto Scaling group scales down, removed EC2 instances are deregistered from the Puppet master as soon as possible.
How can the requirement be met?

A. At instance launch time, use EC2 user data to deploy the AWS CodeDeploy agen
B. Use CodeDeploy to install the Puppet agen
C. When the Auto Scaling group scales out, run a script to register the newly deployed instances to the Puppet maste
D. When the Auto Scaling group scales in, use the EC2 Auto Scaling lifecycle hook to trigger de-registration from the Puppet maste
E. EC2_INSTANCE_TERMINATING
F. Bake the AWS CodeDeploy agent into the base AM
G. When the Auto Scaling group scales out, use CodeDeploy to install the Puppet agent, and execute a script to register the newly deployed instances to the Puppet maste
H. When the Auto Scaling group scales in, use the CodeDeploy ApplicationStop lifecycle hook to run a script to de-register the instance from the Puppet master.
I. At instance launch time, use EC2 user data to deploy the AWS CodeDeploy agen
J. When the Auto Scaling group scales out, use CodeDeploy to install the Puppet agent, and run a script to register the newly deployed instances to the Puppet maste
K. When the Auto Scaling group scales in, use the EC2 user data instance stop script to run a script to de-register the instance from the Puppet master.
L. Bake the AWS Systems Manager agent into the base AM
M. When the Auto Scaling group scales out, use the AWS Systems Manager to install the Puppet agent, and run a script to register the newly deployed instances to the Puppet maste
N. When the Auto Scaling group scales in, use the Systems Manager instance stop lifecycle hook to run a script to de-register the instance from the Puppet master.

Answer: C

NEW QUESTION 164
A company requires that its internally facing web application be nighty available The architecture is made up of one Amazon EC2 web server instance and one NAT instance that provides outbound internet access for updates and accessing public data
Which combination of architecture adjustments should the company implement to achieve high availability? (Select TWO.)

A. Add the NAT instance to an EC2 Auto Scaling group that spans multiple Availability Zones Update the route tables
B. Create additional EC2 instances spanning multiple Availability Zones Add an Application Load Balancer to split the load between them
C. Configure an Application Load Balancer in front of the EC2 instance Configure Amazon CloudWatch alarms to recover the EC2 instance upon host failure
D. Replace the NAT instance with a NAT gateway in each Availability Zone Update the route tables
E. Replace the NAT instance with a NAT gateway that spans multiple Availability Zones Update the route tables

Answer: AD

NEW QUESTION 168
A company is using AWS for an application. The Development team must automate its deployments. The team has set up an AWS CodePipeline to deploy the application to Amazon EC2 instances by using AWS CodeDeploy after it has been built using the AWS CodeBuild service.
The team would like to add automated testing to the pipeline to confirm that the application is healthy before deploying it to the next stage of the pipeline using the same code. The team requires a manual approval action before the application is deployed, even if the test is successful. The testing and approval must be accomplished at the lowest costs, using the simplest management solution.
Which solution will meet these requirements?

A. Add a manual approval action after the last deploy action of the pipelin
B. Use Amazon SNS to inform the team of the stage being triggere
C. Next, add a test action using CodeBuild to do the required test
D. At the end of the pipeline, add a deploy action to deploy the application to the next stage.
E. Add a test action after the last deploy action of the pipelin
F. Configure the action to use CodeBuild to perform the required test
G. If these tests are successful, mark the action as successfu
H. Add a manual approval action that uses Amazon SNS to notify the team, and add a deploy action to deploy the application to the next stage.
I. Create a new pipeline that uses a source action that gets the code from the same repository as the first pipelin
J. Add a deploy action to deploy the code to a test environmen
K. Use a test action using AWS Lambda to test the deploymen
L. Add a manual approval action by using Amazon SNS to notify the team, and add a deploy action to deploy the application to the next stage.
M. Add a test action after the last deployment actio
N. Use a Jenkins server on Amazon EC2 to do the required tests and mark the action as successful if the tests pas
O. Create a manual approval action that uses Amazon SQS to notify the team and add a deploy action to deploy the application to the next stage.

Answer: B

Explanation:
https://docs.aws.amazon.com/codebuild/latest/userguide/sample-build-notifications.html

NEW QUESTION 172
A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps Engineer is tasked with minimizing application response times and improving availability for users in both Regions.
Which combination of actions should be taken to address the latency issues? (Choose three.)

A. Create a new DynamoDB table in the new Region with cross-Region replication enabled.
B. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.
C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.
D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.
E. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.
F. Convert the DynamoDB table to a global table.

Answer: CDF

NEW QUESTION 176
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 178
A retail company has adopted AWS OpsWorks for managing its deployments. In the last three months, the company has discovered that some production instances have been restarting without reason. Upon inspection of the AWS CloudTrail logs, a DevOps Engineer determined that those instances were restarted by OpsWorks. The Engineer now wants automated email notifications whenever OpsWorks restarts an instance when the instance is deemed unhealthy or unable to communicate with the service endpoint.
How can the Engineer meet this requirement?

A. Create a Chef recipe to place a cron to run a custom script within the Amazon EC2 instances that sends an email to the team by using Amazon SES if the OpsWorks agent detects an instance failure.
B. Create an Amazon SNS topic and create a subscription for this topic that contains the destination email addres
C. Create an Amazon CloudWatch rule: specify as a source and specify auto-healing in the initiated_by detail
D. Use the SNS topic as a targe
E. aws.opsworks
F. Create an Amazon SNS topic and create a subscription for this topic that contains the destination email addres
G. Create an Amazon CloudWatch rule specify as a source and specify instance-replacement in the initiated_by detail
H. Use the SNS topic as a targe
I. aws.opsworks
J. Create a subscription for this topic that contains the email addres
K. Enable instance restart notifications within the OpsWorks layer and indicate the destination email address for the notification

Answer: B

NEW QUESTION 183
After presenting a working proof of concept for a new application that uses AWS API Gateway, a Developer must set up a team development environment for the project. Due to a tight timeline, the Developer wants to minimize time spent on infrastructure setup, and would like to reuse the code repository created for the proof of concept. Currently, all source code is stored in AWS CodeCommit.
Company policy mandates having alpha, beta, and production stages with separate Jenkins servers to build code and run tests for every stage. The Development Manager must have the ability to block code propagation between admins at any time. The Security team wants to make sure that users will not be able to modify the environment without permission.
How can this be accomplished?

A. Create API Gateway alpha, beta, and production stage
B. Create a CodeCommit trigger to deploy code to the different stages using an AWS Lambda function.
C. Create API Gateway alpha, beta, and production stage
D. Create an AWS CodePipeline that pulls code from the CodeCommit repositor
E. Create CodePipeline actions to deploy code to the API Gateway stages.
F. Create Jenkins servers for the alpha, beta, and production stages on Amazon EC2 instance
G. Create multiple CodeCommit triggers to deploy code to different stages using an AWS Lambda function.
H. Create an AWS CodePipeline pipeline that pulls code from the CodeCommit repositor
I. Create alpha, beta, and production stages with Jenkins servers on CodePipeline.

Answer: D

NEW QUESTION 188
A company is creating a software solution that executes a specific parallel-processing mechanism. The software can scale to tens of servers in some special scenarios. This solution uses a proprietary library that is license-based, requiring that each individual server have a single, dedicated license installed. The company has 200 licenses and is planning to run 200 server nodes concurrently at most.
The company has requested the following features:
"¢ A mechanism to automate the use of the licenses at scale. "¢ Creation of a dashboard to use in the future to verify which licenses are available at any moment.
What is the MOST effective way to accomplish these requirements?

A. Upload the licenses to a private Amazon S3 bucke
B. Create an AWS CloudFormation template with a Mappings section for the license
C. In the template, create an Auto Scaling group to launch the server
D. In the user data script, acquire an available license from the Mappings sectio
E. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.
F. Upload the licenses to an Amazon DynamoDB tabl
G. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the server
H. In the user data script, acquire an available license from the DynamoDB tabl
I. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.
J. Upload the licenses to a private Amazon S3 bucke
K. Populate an Amazon SQS queue with the list of licenses stored in S3. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the server
L. In the user data script acquire an available license from SQ
M. Create an Auto Scaling lifecycle hook, then use it to put the license back in SQS after the instance is terminated.
N. Upload the licenses to an Amazon DynamoDB tabl
O. Create an AWS CLI script to launch the servers by using the parameter --count, with min:max instances to launc
P. In the user data script, acquire an available license from the DynamoDB tabl
Q. Monitor each instance and, in case of failure, replace the instance, then manually update the DynamoDB table.

Answer: D

NEW QUESTION 190
A company runs a database on a single Amazon EC2 instance in a development environment. The data is stored on separate Amazon EBS volumes that are attached to the EC2 instance. An Amazon Route 53 A record has been created and configured to point to the EC2 instance. The company would like to automate the recovery of the database instance when an instance or Availability Zone (AZ) fails. The company also wants to keep its costs low. The RTO is 4 hours and RPO is 12 hours.
Which solution should a DevOps Engineer implement to meet these requirements?

A. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZ
B. Add a lifecycle hook to the Auto Scaling group and define an Amazon CloudWatch Events rule that is triggered when a lifecycle event occur
C. Have the CloudWatch Events rule invoke an AWS Lambda function to detach or attach the Amazon EBS data volumes from the EC2 instance based on the even
D. Configure the EC2 instance UserData to mount the data volumes (retry on failure with a short delay), then start the database and update the Route 53 record.
E. Run the database on two separate EC2 instances in different AZs with one active and the other as a standb
F. Attach the data volumes to the active instanc
G. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function on EC2 instance terminatio

H. The Lambda function launches a replacement EC2 instanc
I. If the terminated instance was the active node, then the function attaches the data volumes to the standby nod
J. Start the database and update the Route 53 record.
K. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZ
L. Create an AWS Lambda function that is triggered by a scheduled Amazon CloudWatch Events rule every 4 hours to take a snapshot of the data volume and apply a ta
M. Have the instance UserData get the latest snapshot, create a new volume from it, and attach and mount the volum
N. Then start the database and update the Route 53 record.
O. Run the database on two separate EC2 instances in different AZ
P. Configure one of the instances as a master and the other as a standb
Q. Set up replication between the master and standby instance
R. Point the Route 53 record to the maste
S. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function upon the EC2 instance terminatio
T. The Lambda function launches a replacement EC2 instanc
. If the terminated instance was the active node, the function promotes the standby to master and points the Route 53 record to it.

Answer: D

NEW QUESTION 194
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 196
A company hosts parts of a Python-based application using AWS Elastic Beanstalk. An Elastic Beanstalk CLI is being used to create and update the environments. The Operations team detected an increase in requests in one of the Elastic Beanstalk environments that caused downtime overnight. The team noted that the policy used for AWS Auto Scaling is NetworkOut. Based on load testing metrics, the team determined that the application needs to scale CPU utilization to improve the resilience of the environments. The team wants to implement this across all environments automatically.
Following AWS recommendations, how should this automation be implemented?

A. Using ebextensions, place a command within the container_commands key to perform an API call tomodify the scaling metric to CPUUtilization for the Auto Scaling configuratio
B. Use leader_only to execute this command in only the first instance launched within the environment.
C. Using ebextensions, create a custom resource that modifies the AWSEBAutoScalingScaleUpPolicy and AWSEBAutoScalingScaleDownPolicy resources to use CPUUtilization as a metric to scale for the Auto Scaling group.
D. Using ebextensions, configure the option setting MeasureName to CPUUtilization within the aws:autoscaling:trigger namespace.
E. Using ebextensions, place a script within the files key and place it in/opt/elasticbeanstalk/hooks/appdeploy/pre to perform an API call to modify the scaling metric to CPUUtilization for the Auto Scaling configuratio
F. Use leader_only to place this script in only the first instance launched within the environment.

Answer: C

NEW QUESTION 200
You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty. What is the reason for this.

A. You do not have the appropriate permissions to access the logs
B. You do not have your CloudWatch metrics correctly configured
C. ELB Access logs are only available for a maximum of one week.
D. Access logging is an optional feature of Elastic Load Balancing that is disabled by default

Answer: D

Explanation:
Clastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Cach log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues.
Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer. Clastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.
For more information on CLB access logs, please refer to the below document link: from AWS
> http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.
html

NEW QUESTION 202
A DevOps Engineer has several legacy applications that all generate different log formats. The Engineer must standardize the formats before writing them to Amazon S3 for querying and analysis. How can this requirement be met at the LOWEST cost?

A. Have the application send its logs to an Amazon EMR cluster and normalize the logs before sending them to Amazon S3.
B. Have the application send its logs to Amazon QuickSight, then use the Amazon QuickSight SPICE engine to normalize the log
C. Do the analysis directly from Amazon QuickSight.
D. Keep the logs in Amazon S3 and use Amazon Redshift Spectrum to normalize the logs in place.
E. Use Amazon Kinesis Agent on each server to upload the logs and have Amazon Kinesis Data Firehose use an AWS Lambda function to normalize the logs before writing them to Amazon.

Answer: D

NEW QUESTION 203
Two teams are working together on different portions of an architecture and are using AWS CloudFormation to manage their resources. One team administers operating system-level updates and patches, while the other team manages application-level dependencies and updates. The Application team must take the most recent AMI when creating new instances and deploying the application.
What is the MOST scalable method for linking these two teams and processes?

A. The Operating System team uses CloudFormation to create new versions of their AMIs and lists the Amazon Resource names (ARNs) of the AMIs in an encrypted Amazon S3 object as part of the stack output sectio
B. The Application team uses a cross-stack reference to load the encrypted S3 object and obtain the most recent AMI ARNs.
C. The Operating System team uses CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMIs, then places the latest AMI ARNs in an encrypted Amazon S3 object as part of the pipeline outpu
D. The Application team uses a cross-stack reference within their own CloudFormation template to get that S3 object location and obtain the most recent AMI ARNs to use when deploying their application.
E. The Operating System team uses CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMI
F. The team then places the AMI ARNs as parameters in AWS Systems Manager Parameter Store as part of the pipeline outpu
G. The Application team specifies a parameter of type ssm in their CloudFormation stack to obtain the most recent AMI ARN from the Parameter Store.
H. The Operating System team maintains a nested stack that includes both the operating system and Application team template
I. The Operating System team uses a stack update to deploy updates to theapplication stack whenever the Application team changes the application code.

Answer: B

NEW QUESTION 205
A company needs to introduce automatic DNS failover for a distributed web application to a disaster recovery or standby installation. The DevOps Engineer plans to configure Amazon Route 53 to provide DNS routing to alternate endpoint in the event of an application failure.
What steps should the Engineer take to accomplish this? (Select TWO.)

A. Create Amazon Route 53 health checks for each endpoint that cannot be entered as alias record
B. Ensure firewall and routing rules allow Amazon Route 53 to send requests to the endpoints that are specified in the health checks.
C. Create alias records that route traffic to AWS resources and set the value of the Evaluate Target Health option to Yes, then create all the non-alias records.
D. Create a governing Amazon Route 53 record set, set it to failover, and associate it with the primary and secondary Amazon Route 53 record sets to distribute traffic to healthy DNS entries.
E. Create an Amazon CloudWatch alarm to monitor the primary Amazon Route 53 DNS entr
F. Then create an associated AWS Lambda function to execute the failover API call to Route 53 to the secondary DNS entry.
G. Map the primary and secondary Amazon Route 53 record sets to an Amazon CloudFront distribution using primary and secondary origins.

Answer: AC

NEW QUESTION 210
......