AWS-Certified-DevOps-Engineer-Professional Dumps

Amazon AWS Certified DevOps Engineer Professional

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

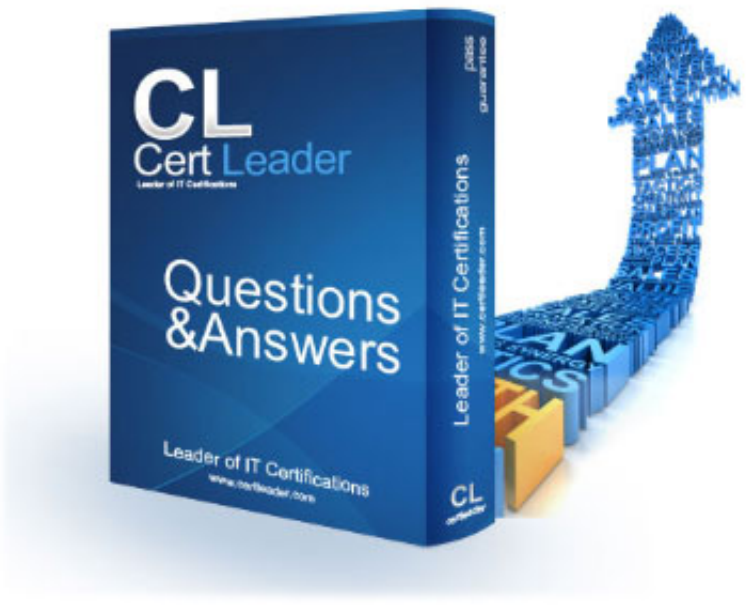

NEW QUESTION 1
A DevOps Engineer must track the health of a stateless RESTful service sitting behind a Classic Load Balancer. The deployment of new application revisions is through a CI/CD pipeline. If the service's latency increases beyond a defined threshold, deployment should be stopped until the service has recovered.
Which of the following methods allow for the QUICKEST detection time?

A. Use Amazon CloudWatch metrics provided by Elastic Load Balancing to calculate average latency.Alarm and stop deployment when latency increases beyond the defined threshold.
B. Use AWS Lambda and Elastic Load Balancing access logs to detect average latenc
C. Alarm and stop deployment when latency increases beyond the defined threshold.
D. Use AWS CodeDeploy's Minimum Healthy Hosts setting to define thresholds for rolling back deployment
E. If these thresholds are breached, roll back the deployment.
F. Use Metric Filters to parse application logs in Amazon CloudWatch Log
G. Create a filter for latency.Alarm and stop deployment when latency increases beyond the defined threshold.

Answer: A

Explanation:
https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-cloudwatch-metrics.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployments-stop.html

NEW QUESTION 2
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 3
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 4
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 5
An n-tier application requires a table in an Amazon RDS MySQL DB instance to be dropped and repopulated at each deployment. This process can take several minutes and the web tier cannot come online until the process is complete. Currently, the web tier is configured in an Amazon EC2 Auto Scaling group, with instances being terminated and replaced at each deployment. The MySQL table
is populated by running a SQL query through an AWS CodeBuild job.
What should be done to ensure that the web tier does not come online before the database is completely configured?

A. Use Amazon Aurora as a drop-in replacement for RDS MySQ
B. Use snapshots to populate the table with the correct data.
C. Modify the launch configuration of the Auto Scaling group to pause user data execution for 600 seconds, allowing the table to be populated.
D. Use AWS Step Functions to monitor and maintain the state of data populatio
E. Mark the database in service before continuing with the deployment.
F. Use an EC2 Auto Scaling lifecycle hook to pause the configuration of the web tier until the table is populated.

Answer: D

NEW QUESTION 6
A company is using an AWS CloudFormation template to deploy web applications. The template requires that manual changes be made for each of the three major environments: production, staging, and development. The current sprint includes the new implementation and configuration of AWS CodePipeline for automated deployments.
What changes should the DevOps Engineer make to ensure that the CloudFormation template is reusable across multiple pipelines?

A. Use a CloudFormation custom resource to query the status of the CodePipeline to determine which environment is launche
B. Dynamically alter the launch configuration of the Amazon EC2 instances.
C. Set up a CodePipeline pipeline for each environment to use input parameter
D. Use CloudFormation mappings to switch associated UserData for the Amazon EC2 instances to match the environment being launched.
E. Set up a CodePipeline pipeline that has multiple stages, one for each development environmen
F. Use AWS Lambda functions to trigger CloudFormation deployments to dynamically alter the UserData of the Amazon EC2 instances launched in each environment.
G. Use CloudFormation input parameters to dynamically alter the LaunchConfiguration and UserData sections of each Amazon EC2 instance every time the CloudFormation stack is updated.

Answer: B

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/continuous-delivery-codepipeline-paramet

NEW QUESTION 7
You have deployed an application to AWS which makes use of Autoscaling to launch new instances. You now want to change the instance type for the new instances. Which of the following is one of the action items to achieve this deployment?

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

Answer: C

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

NEW QUESTION 8
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 9
A consulting company was hired to assess security vulnerabilities within a client company's application and propose a plan to remediate all identified issues. The architecture is identified as follows: Amazon S3 storage for content, an Auto Scaling group of Amazon EC2 instances behind an Elastic Load Balancer with attached Amazon EBS storage, and an Amazon RDS MySQL database. There are also several AWS Lambda functions that communicate directly with the RDS database using connection string statements in the code.
The consultants identified the top security threat as follows: the application is not meeting its requirement to have encryption at rest.
What solution will address this issue with the LEAST operational overhead and will provide monitoring for
potential future violations?

A. Enable SSE encryption on the S3 buckets and RDS databas
B. Enable OS-based encryption of data on EBS volume
C. Configure Amazon Inspector agents on EC2 instances to report on insecure encryption cipher
D. Set up AWS Config rules to periodically check for non-encrypted S3 objects.
E. Configure the application to encrypt each file prior to storing on Amazon S3. Enable OS-based encryption of data on EBS volume
F. Encrypt data on write to RD
G. Run cron jobs on each instance to check for encrypted data and notify via Amazon SN
H. Use S3 Events to call an AWS Lambda function and verify if the file is encrypted.
I. Enable Secure Sockets Layer (SSL) on the load balancer, ensure that AWS Lambda is using SSL to communicate to the RDS database, and enable S3 encryptio
J. Configure the application to force SSL for incoming connections and configure RDS to only grant access if the session is encrypte
K. Configure Amazon Inspector agents on EC2 instances to report on insecure encryption ciphers.
L. Enable SSE encryption on the S3 buckets, EBS volumes, and the RDS databas
M. Store RDS credentials in EC2 Parameter Stor
N. Enable a policy on the S3 bucket to deny unencrypted put
O. Set up AWS Config rules to periodically check for non-encrypted S3 objects and EBS volumes, and to ensure that RDS storage is encrypted.

Answer: D

NEW QUESTION 10
A DevOps Engineer must create a Linux AMI in an automated fashion. The newly created AMI identification must be stored in a location where other build pipelines can access the new identification programmatically
What is the MOST cost-effective way to do this?

A. Build a pipeline in AWS CodePipeline to download and save the latest operating system Open Virtualization Format (OVF) image to an Amazon S3 bucket, then customize the image using the guestfish utilit
B. Use the virtual machine (VM) import command to convert the OVF to an AMI, and store the AMI identification output as an AWS Systems Manager parameter.
C. Create an AWS Systems Manager automation document with values instructing how the image should be create
D. Then build a pipeline in AWS CodePipeline to execute the automation document to build the AMI when triggere
E. Store the AMI identification output as a Systems Manager parameter.
F. Build a pipeline in AWS CodePipeline to take a snapshot of an Amazon EC2 instance running the latest version of the applicatio
G. Then start a new EC2 instance from the snapshot and update the running instance using an AWS Lambda functio
H. Take a snapshot of the updated instance, then convert it to an AM
I. Store the AMI identification output in an Amazon DynamoDB table.
J. Launch an Amazon EC2 instance and install Packe
K. Then configure a Packer build with values defining how the image should be create
L. Build a Jenkins pipeline to invoke the Packer build when triggered to build an AM
M. Store the AMI identification output in an Amazon DynamoDB table.

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 11

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 15
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 18
A company has an application that has predictable peak traffic times. The company wants the application instances to scale up only during the peak times. The application stores state in Amazon DynamoDB. The application environment uses a standard Node.js application stack and custom Chef recipes stored in a private Git repository.
Which solution is MOST cost-effective and requires the LEAST amount of management overhead when performing rolling updates of the application environment?

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

Answer: D

NEW QUESTION 22
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 23
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 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 32
A DevOps Engineer encountered the following error when attempting to use an AWS CloudFormation template to create an Amazon ECS cluster:
An error occurred (InsufficientCapabilitiesException) when calling the CreateStack operation.
What caused this error and what steps need to be taken to allow the Engineer to successfully execute the AWS CloudFormation template?

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

Answer: C

NEW QUESTION 34
A DevOps engineer is tasked with creating a more stable deployment solution for a web application in AWS. Previous deployments have resulted in user-facing bugs, premature user traffic, and inconsistencies between web servers running behind an Application Load Balancer. The current strategy uses AWS CodeCommit to store the code for the application. When developers push to the master branch of the repository. CodeCommit triggers an AWS Lambda deploy function, which invokes an AWS Systems Manager run command to build and deploy the new code to all Amazon EC2 instances.
Which combination of actions should be taken to implement a more stable deployment solution? (Select TWO.)

A. Create a pipeline in AWS CodePipeline with CodeCommit as a source provide
B. Create parallel pipeline stages to build and test the applicatio
C. Pass the build artifact to AWS CodeDeploy.
D. Create a pipeline in AWS CodePipeline with CodeCommit as a source provide
E. Create separate pipeline stages to build and then test the applicatio
F. Pass the build artifact to AWS CodeDeploy.
G. Create and use an AWS CodeDeploy application and deployment group to deploy code updates to the EC2 flee
H. Select the Application Load Balancer for the deployment group.
I. Create individual Lambda functions to run all build, test, and deploy actions using AWS CodeDeploy instead of AWS Systems Manager.
J. Modify the Lambda function to build a single application package to be shared by all instance
K. Use AWS CodeDeploy instead of AWS Systems Manager to update the code on the EC2 fleet.

Answer: CE

NEW QUESTION 35
A company requires its internal business teams to launch resources through pre-approved AWS CloudFormation templates only. The security team requires automated monitoring when resources drift from their expected state.
Which strategy should be used to meet these requirements?

A. Allow users to deploy Cloud Formation stacks using a CloudFormation service role onl
B. Use CloudFormation drift detection to detect when resources have drifted from their expected state.

C. Allow users to deploy CloudFormation stacks using a CloudFormation service role onl
D. Use AWS Config rules to detect when resources have drifted from their expected state.
E. Allow users to deploy CloudFormation stacks using AWS Service Catalog only Enforce the use of a launch constraint Use AWS Config rules to detect when resources have drifted from their expected state.
F. Allow users to deploy CloudFormation stacks using AWS Service Catalog only Enforce the use of a template constraint Use Amazon EventBridge (Amazon CloudWatch Events) notifications to detect when resources have drifted from their expected state.

Answer: B

NEW QUESTION 36
A company plans to stop using Amazon EC2 key pairs for SSH access, and instead plans to use AWS Systems Manager Session Manager. To further enhance security, access to Session Manager must take place over a private network only.
Which combinations of actions will accomplish this? (Select TWO.)

A. Allow inbound access to TCP port 22 in all associated EC2 security groups from the VPC CIDR range.
B. Attach an IAM policy with the necessary Systems Manager permissions to the existing IAM instance profile.
C. Create a VPC endpoint for Systems Manager in the desired Region.
D. Deploy a new EC2 instance that will act as a bastion host to the rest of the EC2 instance fleet.
E. Remove any default routes in the associated route tables.

Answer: BC

NEW QUESTION 41
A DevOps Engineer administers an application that manages video files for a video production company. The application runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. Data is stored in an Amazon RDS PostgreSQL Multi-AZ DB instance, and the video files are stored in an Amazon S3 bucket. On a typical day, 50 GB of new video are added to the S3 bucket. The Engineer must implement a multi-region disaster recovery plan with the least data loss and the lowest recovery times. The current application infrastructure is already described using AWS CloudFormation.
Which deployment option should the Engineer choose to meet the uptime and recovery objectives for the system?

A. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Create an Amazon RDS read replica in the second regio
B. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucke
C. To fail over, promote the read replica as maste
D. Update the CloudFormation stack and increase the capacity of the Auto Scaling group.
E. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Create a scheduled task to take daily Amazon RDS cross-region snapshots to the second regio
F. In the second region, enable cross-region replication between the original S3 bucket and Amazon Glacie
G. In a disaster, launch a new application stack in the second region and restore the database from the most recent snapshot.
H. Launch the application from the CloudFormation template in the second region which sets the capacity of the Auto Scaling group to 1. Use Amazon CloudWatch Events to schedule a nightly task to take a snapshot of the database, copy the snapshot to the second region, and replace the DB instance in the second region from the snapsho
I. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucke
J. To fail over, increase the capacity of the Auto Scaling group.
K. Use Amazon CloudWatch Events to schedule a nightly task to take a snapshot of the database and copy the snapshot to the second regio
L. Create an AWS Lambda function that copies each object to a new S3 bucket in the second region in response to S3 event notification
M. In the second region, launch the application from the CloudFormation template and restore the database from the most recent snapshot.

Answer: A

NEW QUESTION 46
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 48
A company is building a solution for storing files containing Personally Identifiable Information (PII) on AWS.
Requirements state:
*All data must be encrypted at rest and in transit.
*All data must be replicated in at least two locations that are at least 500 miles apart. Which solution meets these requirements?

A. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apar
B. Use a bucket policy to enforce access to the buckets only through HTTP

C. Use a bucket policy to enforce Amazon S3 SSE-C on all objects uploaded to the bucke
D. Configure cross-region replication between the two buckets.
E. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apar
F. Use a bucket policy to enforce access to the buckets only through HTTP
G. Use a bucket policy to enforce S3-Managed Keys (SSE-S3) on all objects uploaded to the bucke
H. Configurecross-region replication between the two buckets.
I. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apar
J. Use an IAM role to enforce access to the buckets only through HTTP
K. Use a bucket policy to enforce Amazon S3-Managed Keys (SSE-S3) on all objects uploaded to the bucke
L. Configure crossregion replication between the two buckets.
M. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apar
N. Use a bucket policy to enforce access to the buckets only through HTTP
O. Use a bucket policy to enforce AWS KMS encryption on all objects uploaded to the bucke
P. Configure cross-region replication between the two bucket
Q. Create a KMS Customer Master Key (CMK) in the primary region for encrypting objects.

Answer: B

NEW QUESTION 49
An application is being deployed with two Amazon EC2 Auto Scaling groups, each configured with an Application Load Balancer. The application is deployed to one of the Auto Scaling groups and an Amazon Route 53 alias record is pointed to the Application Load Balancer of the last deployed Auto Scaling group. Deployments alternate between the two Auto Scaling groups.
Home security devices are making requests into the application. The Development team notes that new requests are coming into the old stack days after the deployment. The issue is caused by devices that are not observing the Time to Live (TTL) setting on the Amazon Route 53 alias record.
What steps should the DevOps Engineer take to address the issue with requests coming to the old stacks, while creating minimal additional resources?

A. Create a fleet of Amazon EC2 instances running HAProxy behind an Application Load Balance
B. The HAProxy instances will proxy the requests to one of the existing Auto Scaling group
C. After a deployment the HAProxy instances are updated to send requests to the newly deployed Auto Scaling group.
D. Reduce the application to one Application Load Balance
E. Create two target groups named Blue and Gree
F. Create a rule on the Application Load Balancer pointed to a single target grou
G. Add logic to the deployment to update the Application Load Balancer rule to the target group of the newly deployed Auto Scaling group.
H. Move the application to an AWS Elastic Beanstalk application with two environment
I. Perform new deployments on the non-live environmen
J. After a deployment, perform an Elastic Beanstalk CNAME swap to make the newly deployed environment the live environment.
K. Create an Amazon CloudFront distributio
L. Set the two existing Application Load Balancers as origins on the distributio
M. After a deployment, update the CloudFront distribution behavior to send requests to the newly deployed Auto Scaling group.

Answer: B

NEW QUESTION 50
A company is migrating an application to AWS that runs on a single Amazon EC2 instance. Because of licensing limitations, the application does not support horizontal scaling. The application will be using Amazon Aurora for its database.
How can the DevOps Engineer architect automated healing to automatically recover from EC2 and Aurora failures, in addition to recovering across Availability Zones (AZs), in the MOST cost-effective manner?

A. Create an EC2 Auto Scaling group with a minimum and maximum instance count of 1, and have it spanacross AZ
B. Use a single-node Aurora instance.
C. Create an EC2 instance and enable instance recover
D. Create an Aurora database with a read replica in a second AZ, and promote it to a primary database instance if the primary database instance fails.
E. Create an Amazon CloudWatch Events rule to trigger an AWS Lambda function to start a new EC2 instance in an available AZ when the instance status reaches a failure stat
F. Create an Aurora database with a read replica in a second AZ, and promote it to a primary database instance when the primary database instance fails.
G. Assign an Elastic IP address on the instanc
H. Create a second EC2 instance in a second A
I. Create an Amazon CloudWatch Events rule to trigger an AWS Lambda function to move the Elastic IP address to the second instance when the first instance fail
J. Use a single-node Aurora instance.

Answer: C

NEW QUESTION 52
A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline. A DevOps Engineer has notices there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation, the DevOps Engineer believes the failures are due to database changes the CloudFormation stack for the application lambda function begins executing. How should the DevOps Engineer overcome this?

A. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before traffic can flow to the new version of the Lambda function
B. Add an AfterAllowTraffic hook to the AppSpec file that forces traffic to wait for any pending database changes before allowing the new version of the Lambda function to respond
C. Add a BeforeInstall hook to the AppSpec file that tests and waits for any necessary database changes before deploying the new version of the Lambda function
D. Add a ValidateService hook to the AppSpec file that inspects incoming traffic and rejects the payload if dependent services such as the database are not yet ready

Answer: B

NEW QUESTION 57

A company uses federated access for its AWS environment The available roles are created and managed using AWS CloudFormation from a CI/CD pipeline. All changes should be made to the IAM roles through the pipeline. The security team found that changes are being made to the roles out-of-band and would like to detect when this occurs.
Which action will accomplish this?

A. Use Amazon Inspector rules to detect and notify when a CloudFormation stack has a configuration change.
B. Use an AWS Trusted Advisor CloudWatch Events rule to detect and notify when a CloudFormation stack has a configuration change.
C. Use AWS CloudTrail to detect and notify when a CloudFormation stack has detected a configuration change.
D. Use an AWS Config rule to detect and notify when a CloudFormation stack has detected a configuration change.

Answer: D

NEW QUESTION 62
A Development team is currently using AWS CodeDeploy to deploy an application revision to an Auto Scaling group. If the deployment process fails, it must be rolled back automatically and a notification must be sent.
What is the MOST effective configuration that can satisfy all of the requirements?

A. Create Amazon CloudWatch Events rules for CodeDeploy operation
B. Configure a CloudWatch Events rule to send out an Amazon SNS message when the deployment fail
C. Configure CodeDeploy to automatically roll back when the deployment fails.
D. Use available Amazon CloudWatch metrics for CodeDeploy to create CloudWatch alarm
E. Configure CloudWatch alarms to send out an Amazon SNS message when the deployment fail
F. Use AWS CLI toredeploy a previously deployed revision.
G. Configure a CodeDeploy agent to create a trigger that will send notification to Amazon SNS topics when the deployment fail
H. Configure CodeDeploy to automatically roll back when the deployment fails.
I. Use AWS CloudTrail to monitor API calls made by or on behalf of CodeDeploy in the AWS account.Send an Amazon SNS message when deployment fail
J. Use AWS CLI to redeploy a previously deployed revision.

Answer: C

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/monitoring-sns-event-notifications-create-trigger.htm

NEW QUESTION 64
Company policies require that information about IP traffic going between instances in the production Amazon VPC is captured. The capturing mechanism must always be enabled and the Security team must be notified when any changes in configuration occur.
What should be done to ensure that these requirements are met?

A. Using the UserData section of an AWS CloudFormation template, install tcpdump on every provisioned Amazon EC2 instanc
B. The output of the tool is sent to Amazon EFS for aggregation and queryin
C. In addition, scheduling an Amazon CloudWatch Events rule calls an AWS Lambda function to check whether tcpdump is up and running and sends an email to the security organization when there is an exception.
D. Create a flow log for the production VPC and assign an Amazon S3 bucket as a destination for delivery.Using Amazon S3 Event Notification, set up an AWS Lambda function that is triggered when a new log file gets delivere
E. This Lambda function updates an entry in Amazon DynamoDB, which is periodically checked by scheduling an Amazon CloudWatch Events rule to notify security when logs have not arrived.
F. Create a flow log for the production VP
G. Create a new rule using AWS Config that is triggered by configuration changes of resources of type '"˜EC2:VPC'. As part of configuring the rule, create an AWS Lambda function that looks up flow logs for a given VP
H. If the VPC flow logs are not configured, return a '"˜NON_COMPLIANT' status and notify the security organization.
I. Configure a new trail using AWS CloudTrail servic
J. Using the UserData section of an AWS CloudFormation template, install tcpdump on every provisioned Amazon EC2 instanc
K. Connect Amazon Athena to the CloudTrail and write an AWS Lambda function that monitors for a flow log disable even
L. Once the CloudTrail entry has been spotted, alert the security organization

Answer: C

NEW QUESTION 69
You are responsible for your company's large multi-tiered Windows-based web application running on Amazon EC2 instances situated behind a load balancer. While reviewing metrics, you've started noticing an upwards trend for slow customer page load time. Your manager has asked you to come up with a solution to ensure that customer load time is not affected by too many requests per second. Which technique would you use to solve this issue?

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

Answer: C

Explanation:
Auto Scaling helps you ensure that you have the correct number of Amazon CC2 instances available to handle the load for your application. You create collections of
CC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Auto Scaling ensures that your group
never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Auto Scaling ensures that yourgroup never goes

above this size. If you specify the desired capacity, either when you create the group or at any time thereafter.
Auto Scaling ensures that yourgroup has this many
instances. If you specify scaling policies, then Auto Scaling can launch or terminate instances as demand on your application increases or decreases.
Option A and B are invalid because Autoscaling is required to solve the issue to ensure the application can handle high traffic loads.
Option D is invalid because there is no Autoscaling template.
For more information on Autoscaling, please refer to the below document link: from AWS

- http://docs.aws.amazon.com/autoscaling/latest/userguide/Whatl sAutoScaling.html

NEW QUESTION 73
A DevOps Engineer is deploying an Amazon API Gateway API with an AWS Lambda function providing the backend functionality. The Engineer needs to record the source IP address and response status of every API call.
Which combination of actions should the DevOps Engineer take to implement this functionality? (Choose three.)

A. Configure AWS X-Ray to enable access logging for the API Gateway requests.
B. Configure the API Gateway stage to enable access logging and choose a logging format.
C. Create a new Amazon CloudWatch Logs log group or choose an existing log group to store the logs.
D. Grant API Gateway permission to read and write logs to Amazon CloudWatch through an IAM role.
E. Create a new Amazon S3 bucket or choose an existing S3 bucket to store the logs.
F. Configure API Gateway to stream its log data to Amazon Kinesis.

Answer: BDE

NEW QUESTION 78
A DevOps Engineer is implementing a mechanism for canary testing an application on AWS. The application was recently modified and went through security, unit, and functional testing. The application needs to be deployed on an AutoScaling group and must use a Classic Load Balancer.
Which design meets the requirement for canary testing?

A. Create a different Classic Load Balancer and Auto Scaling group for blue/green environment
B. Use Amazon Route 53 and create weighted A records on Classic Load Balancer.
C. Create a single Classic Load Balancer and an Auto Scaling group for blue/green environment
D. Use Amazon Route 53 and create A records for Classic Load Balancer IP
E. Adjust traffic using A records.
F. Create a single Classic Load Balancer and an Auto Scaling group for blue/green environment
G. Create an Amazon CloudFront distribution with the Classic Load Balancer as the origi
H. Adjust traffic using CloudFront.
I. Create a different Classic Load Balancer and Auto Scaling group for blue/green environment
J. Create an Amazon API Gateway with a separate stage for the Classic Load Balance
K. Adjust traffic by giving weights to this stage.

Answer: A

NEW QUESTION 82
A law firm is running a web application on AWS. The system manages legal documents uploaded by users, and stores the documents in Amazon S3. Users have complained that file uploads are taking too long and there are timeouts during peak usage. A DevOps engineer found that web servers are managing concurrent uploads and are overloaded.
Which actions should be taken to troubleshoot the issue in the MOST cost-effective manner?

A. Create an AWS CloudFront distribution in front of the web servers, and modify the application to upload to Amazon S3 using S3 Transfer Acceleration.
B. Modify the application so the browser uses a signed URL to directly upload to Amazon S3 using multipart uploads.
C. Create an AWS CloudFront distribution in front of the web servers, and modify the application to store files in Amazon EFS in the Max I/O performance mode.
D. Place the web servers in an Amazon EC2 Auto Scaling group to include Spot Instances and modify the application to upload to Amazon S3 using multipart uploads.

Answer: A

NEW QUESTION 84
A company is using several AWS CloudFormation templates for deploying infrastructure as code. In most of the deployments, the company uses Amazon EC2 Auto Scaling groups. A DevOps Engineer needs to update the AMIs for the Auto Scaling group in the template if newer AMIs are available.
How can these requirements be met?

A. Manage the AMI mappings in the CloudFormation templat
B. Use Amazon CloudWatch Events for detecting new AMIs and updating the mapping in the templat
C. Reference the map in the launch configuration resource block.
D. Use conditions in the AWS CloudFormation template to check if new AMIs are available and return the AMI I
E. Reference the returned AMI ID in the launch configuration resource block.
F. Use an AWS Lambda-backed custom resource in the template to fetch the AMI ID
G. Reference the returned AMI ID in the launch configuration resource block.
H. Launch an Amazon EC2 m4.small instance and run a script on it to check for new AMI
I. If new AMIs are available, the script should update the launch configuration resource block with the new AMI ID.

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/walkthrough-custom-resources-lambda-lo

NEW QUESTION 85
A DevOps Engineer must automate a weekly process of identifying unnecessary permissions on a per-user basis, across all users in an AWS account. This

process should evaluate the permissions currently granted to each user by examining the user's attached IAM access policies compared to the permissions the user has actually used in the past 90 days. Any differences in the comparison would indicate that the user has more permissions than are required. A report of the deltas should be sent to the Information Security team for further review and IAM user access policy revisions, as required.
Which solution is fully automated and will produce the MOST detailed deltas report?

A. Create an AWS Lambda function that calls the IAM Access Advisor API to pull service permissions granted on a user-by-user basis for all users in the AWS accoun
B. Ensure that Access Advisor is configured with a tracking period of 90 day
C. Invoke the Lambda function using an Amazon CloudWatch Events rule on a weekly schedul
D. For each record, by user, by service, if the Access Advisor Last Accesses field indicates a day count instead of "Not accesses in the tracking period," this indicates a delta compared to what is in the user's currently attached access police
E. After Lambda has iterated through all users in the AWS account, configure it to generate a report and send the report using Amazon SES.
F. Configure an AWS CloudTrail trail that spans all AWS Regions and all read/write events, and point this trail to an Amazon S3 bucke
G. Create Amazon Athena table and specify the S3 bucket ARN in the CREATE TABLE quer
H. Create an AWS Lambda function that accesses the Athena table using the SDK, which performs a SELECT, ensuring that the WHERE clause includes userIdentity, eventName, and eventTim
I. Compare the results against the user's currently attached IAM access policies to determine any delta
J. Configure an Amazon CloudWatch Events schedule to automate this process to run once a wee
K. Configure Amazon SES to send a consolidated report to the Information Security team.
L. Configure VPC Flow Logs on all subnets across all VPCs in all regions to capture user traffic across the entire accoun
M. Ensure that all logs are being sent to a centralized Amazon S3 bucket, so all flow logs can be consolidated and aggregate
N. Create an AWS Lambda function that is triggered once a week by an Amazon CloudWatch Events schedul
O. Ensure that the Lambda function parses the flow log files for the following information: IAM user ID, subnet ID, VPC ID, Allow/Reject status per API call, and service nam
P. Then have the function determine the deltas on a user-by-user basi
Q. Configure the Lambda function to send the consolidated report using Amazon SES.
R. Create an Amazon ES cluster and note its endpoint URL, which will be provided as an environment variable into a Lambda functio
S. Configure an Amazon S3 event on a AWS CloudTrail trail destination S3 bucket and ensure that the event is configured to send to a Lambda functio
T. Create the Lambda function to consume the events, parse the input from JSON, and transform it to an Amazon ES document forma
. POST the documents to the Amazon ES cluster's endpoint by way of the passed-in environment variabl
. Make sure that the proper indexing exists in Amazon ES and use Apache Lucene queries to parse the permissions on a user-by-user basi
. Export the deltas into a report and have Amazon ES send the reports to the Information Security team using Amazon SES every week.

Answer: C

NEW QUESTION 86
A company has several AWS accounts. The accounts are shared and used across multiple teams globally, primarily for Amazon EC2 instances. Each EC2 instance has tags for team, environment, and cost center to ensure accurate cost allocations.
How should a DevOps Engineer help the teams audit their costs and automate infrastructure cost optimization across multiple shared environments and accounts?

A. Set up a scheduled script on the EC2 instances to report utilization and store the instances in an Amazon DynamoDB tabl
B. Create a dashboard in Amazon QuickSight with DynamoDB as the source data to find underutilized instance
C. Set up triggers from Amazon QuickSight in AWS Lambda to reduce underutilized instances.
D. Create a separate Amazon CloudWatch dashboard for EC2 instance tags based on cost center, environment, and team, and publish the instance tags out using unique links for each tea
E. For each team, set up a CloudWatch Events rule with the CloudWatch dashboard as the source, and set up a trigger to initiate an AWS Lambda function to reduce underutilized instances.
F. Create an Amazon CloudWatch Events rule with AWS Trusted Advisor as the source for low utilization EC2 instance
G. Trigger an AWS Lambda function that filters out reported data based on tags for each team, environment, and cost center, and store the Lambda function in Amazon S3. Set up a second trigger to initiate a Lambda function to reduce underutilized instances.
H. Use AWS Systems Manager to track instance utilization and report underutilized instances to Amazon CloudWatc
I. Filter data in CloudWatch based on tags for team, environment, and cost cente
J. Set up triggers from CloudWatch into AWS Lambda to reduce underutilized instances

Answer: C

Explanation:
https://github.com/aws/Trusted-Advisor-Tools/tree/master/LowUtilizationEC2Instances https://docs.aws.amazon.com/quicksight/latest/user/supported-data-sources.html

NEW QUESTION 91
A company is using Docker containers for an application deployment and wants to move its application to AWS. The company currently manages its own clusters on premises to manage the deployment of these containers. It wants to deploy its application to a managed service in AWS and wants the entire flow of the deployment process to be automated. In addition, the company has the following requirements:
Focus first on the development workload. The environment must be easy to manage.
Deployment should be repeatable and reusable for new environments. Store the code in a GitHub repository.
Which solution will meet these requirements?

A. Set up an Amazon ECS environmen
B. Use AWS CodePipeline to create a pipeline that is triggered on a commit to the GitHub repositor
C. Use AWS CodeBuild to create the container images and AWS CodeDeploy to publish the container image to the ECS environment.
D. Use AWS CodePipeline that triggers on a commit from the GitHub repository, build the container images with AWS CodeBuild, and publish the container images to Amazon EC
E. In the final stage, use AWS CloudFormation to create an Amazon ECS environment that gets the container images from the ECR repository.
F. Create a Kubernetes Cluster on Amazon EC2. Use AWS CodePipeline to create a pipeline that is triggered when the code is committed to the repositor
G. Create the container images with a Jenkins server on EC2 and store them in the Docker Hu
H. Use AWS Lambda from the pipeline to trigger the deployment to the Kubernetes Cluster.
I. Set up an Amazon ECS environmen
J. Use AWS CodePipeline to create a pipeline that is triggered on a commit to the GitHub repositor
K. Use AWS CodeBuild to create the container and store it in the Docker Hu
L. Use an AWS Lambda function to trigger a deployment and pull the new container image from the Docker Hub.

Answer: A

NEW QUESTION 95
An ecommerce company uses a large number of Amazon EBS backed Amazon EC2 instances. To decrease manual work across all the instances, a DevOps engineer is tasked with automating restart actions when EC2 instance retirement events are scheduled.
How can this be accomplished?

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

Answer: D

NEW QUESTION 99
A Developer is maintaining a fleet of 50 Amazon EC2 Linux servers. The servers are part of an Amazon EC2 Auto Scaling group, and also use Elastic Load Balancing for load balancing.
Occasionally, some application servers are being terminated after failing ELB HTTP health checks. The Developer would like to perform a root cause analysis on the issue, but before being able to access application logs, the server is terminated.
How can log collection be automated?

A. Use Auto Scaling lifecycle hooks to put instances in a Pending:Wait stat
B. Create an Amazon CloudWatch Alarm for EC2 Instance Terminate and trigger an AWS Lambda function that executes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the Successful lifecycle action once logs are collected.
C. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
D. Create a Config rule for EC2 Instance-terminate Lifecycle and trigger a step function that executes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collecte
E. Action
F. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
G. Create an Amazon CloudWatch subscription filter for EC2 Instance and trigger a CloudWatch agent that executes a script to called logs, push them to Amazon S3, and complete the lifecycle action Terminate Successful once logs are collected.
H. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
I. Create an Amazon CloudWatch Events rule for EC2 Instance- and trigger an AWS Lambda function that executes a SSM Run Command script to collect logs, push them to Amazon S3, terminate Lifecycle Action and complete the lifecycle action once logs are collected.

Answer: D

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html

NEW QUESTION 100
A web application with multiple services runs on Amazon EC2 instances behind an Application Load Balancer. The application stores data in an Amazon RDS Multi-AZ DB instance. The instance health check used by the load balancer returns PASS if at least one service is running on the instance.
The company uses AWS CodePipeline with AWS CodeBuild and AWS CodeDeploy steps to deploy code to test and production environments. Recently, a new version was unable to connect to the database server in the test environment. One process was running, so the health checks reported healthy and the application was promoted to production, causing a production outage. The company wants to ensure that test builds are fully functional before a promotion to production.
Which changes should a DevOps Engineer make to the test and deployment process? (Choose two.)

A. Add an automated functional test to the pipeline that ensures solid test cases are performed.
B. Add a manual approval action to the CodeDeploy deployment pipeline that requires a Testing Engineer to validate the testing environment.
C. Refactor the health check endpoint the Elastic Load Balancer is checking to better validate actual application functionality.
D. Refactor the health check endpoint the Elastic Load Balancer is checking to return a text-based status result and configure the load balancer to check for a valid response.
E. Add a dependency checking step to the existing testing framework to ensure compatibility.

Answer: DE

NEW QUESTION 101
A DevOps Engineer must implement monitoring for a workload running on Amazon EC2 and Amazon RDS MySQL. The monitoring must include:
Application logs and operating system metrics for the Amazon EC2 instances Database logs and operating system metrics for the Amazon RDS database Which steps should the Engineer take?

A. Install an Amazon CloudWatch agent on the EC2 and RDS instance
B. Configure the agent to send the operating system metrics and application and database logs to CloudWatch.
C. Install an Amazon CloudWatch agent on the EC2 instance, and configure the agent to send the application logs and operating system metrics to CloudWatc
D. Enable RDS Enhanced Monitoring, and modify the RDS instance to publish database logs to CloudWatch Logs.
E. Install an Amazon CloudWatch Logs agent on the EC2 instance and configure it to send application logs to CloudWatch.
F. Set up scheduled tasks on the EC2 and RDS instances to put operating system metrics and applicationand database logs into an Amazon S3 bucke
G. Set up an event on the bucket to invoke an AWS Lambda function to monitor for errors each time an object is put into the bucket.

Answer: B

NEW QUESTION 104

The Security team depends on AWS CloudTrail to detect sensitive security issues in the company's AWS account. The DevOps Engineer needs a solution to auto-remediate CloudTrail being turned off in an AWS account.
What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries?

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

Answer: A

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

NEW QUESTION 109
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 114
You have an application running a specific process that is critical to the application's functionality, and have added the health check process to your Auto Scaling Group. The instances are showing healthy but the application itself is not working as it should. What could be the issue with the health check, since it is still showing the instances as healthy.

A. You do not have the time range in the health check properly configured
B. It is not possible for a health check to monitor a process that involves the application
C. The health check is not configured properly
D. The health check is not checking the application process

Answer: D

Explanation:
If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance
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 119
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 120
A DevOps Engineer is developing a deployment strategy that will allow for data-driven decisions before a feature is fully approved for general availability. The current deployment process uses AWS CloudFormation and blue/green-style deployments. The development team has decided that customers should be randomly assigned to groups, rather than using a set percentage, and redirects should be avoided.
What process should be followed to implement the new deployment strategy?

A. Configure Amazon Route 53 weighted records for the blue and green stacks, with 50% of trafficconfigured to route to each stack.
B. Configure Amazon CloudFront with an AWS Lambda@Edge function to set a cookie when CloudFront receives a reques
C. Assign the user to a version A or B, and configure the web server to redirect to version A or B.
D. Configure Amazon CloudFront with an AWS Lambda@Edge function to set a cookie when CloudFront receives a reques
E. Assign the user to a version A or B, then return the corresponding version to the viewer.
F. Configure Amazon Route 53 with an AWS Lambda function to set a cookie when Amazon CloudFront receives a reques
G. Assign the user to version A or B, then return the corresponding version to the viewer.

Answer: C

Explanation:
https://docs.aws.amazon.com/zh_cn/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html

NEW QUESTION 125
A company uses AWS KMS with CMKs and manual key rotation to meet regulatory compliance requirements. The security team wants to be notified when any keys have not been rotated after 90 days.
Which solution will accomplish this?

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

Answer: C

NEW QUESTION 128
A company has developed a Node.js web application which provides REST services to store and retrieve time series data. The web application is built by the Development team on company laptops, tested locally, and manually deployed to a single on-premises server, which accesses a local MySQL database. The company is starting a trial in two weeks, during which the application will undergo frequent updates based on customer feedback. The following requirements must be met:
*The team must be able to reliably build, test, and deploy new updates on a daily basis, without downtime or degraded performance.
*The application must be able to scale to meet an unpredictable number of concurrent users during the trial. Which action will allow the team to quickly meet these objectives?

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

Answer: C

NEW QUESTION 129
A company is using AWS Organizations and wants to implement a governance strategy with the following requirements:
• AWS resource access is restricted to the same two Regions for all accounts.
• AWS services are limited to a specific group of authorized services for all accounts.
• Authentication is provided by Active Directory.
• Access permissions are organized by job function and are identical in each account. Which solution will meet these requirements?

A. Establish an organizational unit (OU) with group policies in the master account to restrict Regions and authorized service
B. Use AWS Cloud Formation StackSets to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.
C. Establish a permission boundary in the master account to restrict Regions and authorized service
D. Use AWS CloudFormation StackSet to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.
E. Establish a service control policy in the master account to restrict Regions and authorized service
F. Use AWS Resource Access Manager to share master account roles with permissions for each job function, including AWS SSO for authentication in each account.
G. Establish a service control policy in the master account to restrict Regions and authorized service
H. Use CloudFormation StackSet to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.

Answer: D

NEW QUESTION 132
A rapidly growing company wants to scale for Developer demand for AWS development environments.

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

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

Answer: A

NEW QUESTION 136
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 138
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 logs Do the analysis directly from Amazon QuickSight.
C. Keep the logs in Amazon S3 and use Amazon Redshift Spectrum to normalize the logs in place
D. 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 S3

Answer: D

NEW QUESTION 141
A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.
Which solution will accomplish this with the LEAST amount of development effort?

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

Answer: B

NEW QUESTION 144
You have an application which consists of EC2 instances in an Auto Scaling group. Between a particular time frame every day, there is an increase in traffic to your website. Hence users are complaining of a poor response time on the application. You have configured your Auto Scaling group to deploy one new EC2 instance when CPU utilization is greater than 60% for 2 consecutive periods of 5 minutes. What is the least cost-effective way to resolve this problem?

A. Decrease the consecutive number of collection periods
B. Increase the minimum number of instances in the Auto Scaling group
C. Decrease the collection period to ten minutes
D. Decrease the threshold CPU utilization percentage at which to deploy a new instance

Answer: B

Explanation:
If you increase the minimum number of instances, then they will be running even though the load is not high on the website. Hence you are incurring cost even though there is no need.
All of the remaining options are possible options which can be used to increase the number of instances on a high load.
For more information on On-demand scaling, please refer to the below link:
> http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html
Note: The tricky part where the question is asking for 'least cost effective way". You got the design consideration correctly but need to be careful on how the question is phrased.

NEW QUESTION 148
A company's legacy application uses IAM user credentials to access resources in the company's AWS Organizations organization. A DevOps engineer needs to ensure new IAM users cannot be created unless the employee creating the IAM user is on an exception list.
Which solution will meet these requirements?

A. Attach an Organizations SCP with an explicit deny for all iam:CreateAccessKey actions with a condition that excludes StringNotEquals for aws:username with a value of the exception list.
B. Attach an Organizations SCP with an explicit deny for all iam:CreateUser actions with a condition that includes StringEquals for aws:username with a value of the exception list.
C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a pattern that matches the iam:CreateAccessKey action with an AWS Lambda function targe
D. The function will check the user name account against an exception lis
E. If the user is not in the exception list, the function will delete the user.
F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a pattern that matches the iam:CreateUser action with an AWS Lambda function targe
G. The function will check the user name and account against an exception lis
H. If the user is not in the exception list, the function will delete the user.

Answer: B

NEW QUESTION 150
A company wants to use AWS Systems Manager documents to bootstrap physical laptops for developers. The bootstrap code is stored in GitHub. A DevOps engineer has already created a Systems Manager activation, installed the Systems Manager agent with the registration code, and installed an activation ID on all the laptops.
Which set of steps should be taken next?

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

Answer: D

NEW QUESTION 155
A media customer has several thousand amazon EC2 instances in an AWS account. The customer is using a Slack channel for team communications and important updates. A DevOps Engineer was told to send all AWS-scheduled EC2 maintenance notifications to the company Slack channel.
Which method should the Engineer use to implement this process in the LEAST amount of steps?

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

Answer: B

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

NEW QUESTION 157
A DevOps Engineer is building a continuous deployment pipeline for a serverless application using AWS CodePipeline and AWS CodeBuild. The source, build, and test stages have been created with the deploy stage remaining. The company wants to reduce the risk of an unsuccessful deployment by deploying to a specified subset of customers and monitoring prior to a full release to all customers.
How should the deploy stage be configured to meet these requirements?

A. Use AWS CloudFormation to publish a new version on every stack updat
B. Then set up a CodePipeline approval action for a Developer to test and approve the new versio
C. Finally, use a CodePipeline invoke action to update an AWS Lambda function to use the production alias
D. Use CodeBuild to use the AWS CLI to update the AWS Lambda function code, then publish a new version of the function and update the production alias to point to the new version of the function.
E. Use AWS CloudFormation to define the serverless application and AWS CodeDeploy to deploy the AWS Lambda functions using DeploymentPreference: . Canary10Percent15Minutes
F. Use AWS CloudFormation to publish a new version on every stack updat
G. Use the RoutingConfig property of the AWS::Lambda::Alias resource to update the traffic routing during the stack update.

Answer: C

Explanation:
https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverle

NEW QUESTION 160
A company runs a three-tier web application in its production environment, which is built on a single AWS CloudFormation template made up of Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. Data is stored in an Amazon RDS Multi-AZ DB instance with read replicas. Amazon Route 53 manages the application's public DNS record.
A DevOps Engineer must create a workflow to mitigate a failed software deployment by rolling back changes in the production environment when a software cutover occurs for new application software.
What steps should the Engineer perform to meet these requirements with the LEAST amount of downtime?

A. Use CloudFormation to deploy an additional staging environment and configure the Route 53 DNS withweighted record
B. During cutover, change the Route 53 A record weights to achieve an even traffic distribution between the two environment
C. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
D. Use a single AWS Elastic Beanstalk environment to deploy the staging and production environments.Update the environment by uploading the ZIP file with the new application cod
E. Swap the Elastic Beanstalk environment CNAM
F. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
G. Use a single AWS Elastic Beanstalk environment and an AWS OpsWorks environment to deploy the staging and production environment
H. Update the environment by uploading the ZIP file with the new application code into the Elastic Beanstalk environment deployed with the OpsWorks stac
I. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
J. Use AWS CloudFormation to deploy an additional staging environment, and configure the Route 53 DNS with weighted record
K. During cutover, increase the weight distribution to have more traffic directed to the new staging environment as workloads are successfully validate
L. Keep the old production environment in place until the new staging environment handles all traffic.

Answer: D

NEW QUESTION 163
A DevOps Engineer needs to design and implement a backup mechanism for Amazon EFS. The Engineer is given the following requirements:
*The backup should run on schedule.
*The backup should be stopped if the backup window expires.
*The backup should be stopped if the backup completes before the backup window.
*The backup logs should be retained for further analysis.
The design should support highly available and fault-tolerant paradigms.
*Administrators should be notified with backup metadata. Which design will meet these requirements?

A. Use AWS Lambda with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activit
B. Run backup scripts on Amazon EC2 in an Auto Scaling grou
C. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading backup logs to Amazon S3. Use Amazon SNS to notify administrators with backup activity metadata.
D. Use Amazon SWF with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activit
E. Run backup scripts on Amazon EC2 in an Auto Scaling grou
F. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading backup logs to Amazon Redshif
G. Use CloudWatch Alarms to notify administrators with backup activity metadata.
H. Use AWS Data Pipeline with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activit
I. Run backup scripts on Amazon EC2 in a single Availability Zon
J. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading the backup logs to Amazon RD
K. Use Amazon SNS to notify administrators with backup activity metadata.
L. Use AWS CodePipeline with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activit
M. Run backup scripts on Amazon EC2 in a single Availability Zon
N. Use Auto Scaling lifecycle hooks and the SSM Run Command on Amazon EC2 for uploading backup logs to Amazon S3. Use Amazon SES to notify admins with backup activity metadata.

Answer: A

Explanation:
https://docs.aws.amazon.com/efs/latest/ug/alternative-efs-backup.html

NEW QUESTION 168
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 172
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 177
A company wants to implement a CI/CD pipeline for an application that is deployed on AWS. The company also has a source-code analysis tool hosted on premises that checks for security flaws. The tool has not yet been migrated to AWS and can be accessed only on premises. The company wants to run checks against the source code as part of the pipeline before the code is compiled. The checks take anywhere from minutes to an hour to complete.
How can a DevOps Engineer meet these requirements?

A. Use AWS CodePipeline to create a pipelin
B. Add an action to the pipeline to invoke an AWS Lambda function after the source stag
C. Have the Lambda function invoke the source-code analysis tool on premises against the source input from CodePipelin
D. The function then waits for the execution to complete and places the output in a specified Amazon S3 location.
E. Use AWS CodePipeline to create a pipeline, then create a custom action typ
F. Create a job worker for the custom action that runs on hardware hosted on premise
G. The job worker handles running security checks with the on-premises code analysis tool and then returns the job results to CodePipelin
H. Have the pipeline invoke the custom action after the source stage.
I. Use AWS CodePipeline to create a pipelin
J. Add a step after the source stage to make an HTTPS request to the on-premises hosted web service that invokes a test with the source code analysis too
K. When the analysis is complete, the web service sends the results back by putting the results in an Amazon S3 output location provided by CodePipeline.
L. Use AWS CodePipeline to create a pipelin
M. Create a shell script that copies the input source code to a location on premise
N. Invoke the source code analysis tool and return the results to CodePipelin
O. Invoke the shell script by adding a custom script action after the source stage.

Answer: B

NEW QUESTION 179
A company uses AWS Storage Gateway in file gateway mode in front of an Amazon S3 bucket that is used by multiple resources. In the morning when business begins, users do not see the objects processed by a third party the previous evening. When a DevOps engineer looks directly at the S3 bucket, the data is there, but it is missing in Storage Gateway.
Which solution ensures that all the updated third-party files are available in the morning?

A. Configure a nightly Amazon EventBridge (Amazon CloudWatch Events) event to trigger an AWS Lambda function to run the RefreshCache command for Storage Gateway.
B. Instruct the third party to put data into the S3 bucket using AWS Transfer for SFTP.
C. Modify Storage Gateway to run in volume gateway mode.
D. Use S3 same-Region replication to replicate any changes made directly in the S3 bucket to Storage Gateway.

Answer: A

NEW QUESTION 182
For auditing, analytics, and troubleshooting purposes, a DevOps Engineer for a data analytics application needs to collect all of the application and Linux system logs from the Amazon EC2 instances before termination. The company, on average, runs 10,000 instances in an Auto Scaling group. The company requires the ability to quickly find logs based on instance IDs and date ranges.
Which is the MOST cost-effective solution?

A. Create an EC2 Instance-terminate Lifecycle Action on the group, write a termination script for pushing logs into Amazon S3, and trigger an AWS Lambda function based on S3 PUT to create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
B. Create an EC2 Instance-terminate Lifecycle Action on the group, write a termination script for pushing logs into Amazon CloudWatch Logs, create a CloudWatch Events rule to trigger an AWS Lambda function to create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
C. Create an EC2 Instance-terminate Lifecycle Action on the group, create an Amazon CloudWatch Events rule based on it to trigger an AWS Lambda function for storing the logs in Amazon S3, and create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
D. Create an EC2 Instance-terminate Lifecycle Action on the group, push the logs into Amazon Kinesis Data Firehouse, and select Amazon ES as the destination for providing storage and search capability.

Answer: C

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html

NEW QUESTION 185
A startup company is developing a web application on AWS. It plans to use Amazon RDS for persistence and deploy the application to Amazon EC2 with an Auto Scaling group. The company would also like to separate the environments for development, testing, and production.
What is the MOST secure and flexible approach to manage the application configuration?

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

Answer: D

NEW QUESTION 189
A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times An application uses a MySQL database and Amazon EC2 web servers. The development learn needs a strategy for failover and disaster recovery
Which combination of deployment strategies will meet these requirements? {Select TWO)

A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store Use Aurora's automatic recovery capabilities in the event of a discluster.
B. Create an Amazon Aurora global database m two Regions as the data store In the event of a failure, promote the secondary Region as the master for the application
C. Create an Amazon Aurora multi-master cluster across multiple Regions as the data store Use an Network Load Balancer to balance the database traffic in different Regions.
D. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Regions Use health checks to determine the availability in a given Regio
E. Use Auto Scaling groups in each Region to adjust capacity based on demand
F. Set up the application m two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on demand in the event of a disaster, adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

Answer: BE

NEW QUESTION 193
A company is developing a web application's infrastructure using AWS CloudFormation. The database engineering team maintains the database resources in a CloudFormation template, and the software development team maintains the web application resources in a separate CloudFormation template As the scope of the application grows, the software development team needs to use resources maintained by the database engineering team However, both teams have their own review and lifecycle management processes that they want to keep Both teams also require resource-level change-set reviews The software development team would like to deploy changes to this template using their CI/CD pipeline.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 194
A company needs to implement a robust CI/CD pipeline to automate the deployment of an application in AWS. The pipeline must support continuous integration, continuous delivery, and automatic rollback upon deployment failure. The entire CI/CD pipeline must be capable of being re-provisioned in alternate AWS accounts or Regions within minutes. A DevOps engineer has already created an AWS CodeCommit repository to store the source code.
Which combination of actions should be taken when building this pipeline to meet these requirements? (Select THREE.)

A. Configure an AWS CodePipehne pipeline with a build stage using AWS CodeBuild.
B. Copy the build artifact from CodeCommit to Amazon S3.
C. Create an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer (ALB) and set the ALB as the deployment target in AWS CodePipeline.
D. Create an AWS Elastic Beanstalk environment as the deployment target in AWS CodePipeline.
E. Implement an Amazon SQS queue to decouple the pipeline components.
F. Provision all resources using AWS CloudFormation.

Answer: ABD

NEW QUESTION 197
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 200
The resources for a business-critical, three-tier web application are expressed in a series of AWS CloudFormation templates. The application is using Amazon RDS for data and Amazon ElastiCache for session state. Users have reported degraded performance in the application. A DevOps Engineer notices that the T2 instance type is being used for the application tier and CPU usage is at 100% in Amazon CloudWatch. What process should the Engineer follow to restore operations with the LEAST amount of distribution to the end users?

A. Write a new CloudFormation template to include Amazon CloudFront in the environment, launch the stack, and update the Amazon Route 53 A record
B. Launch a new CloudFormation stack for the application tier using the M4 instance type, run acceptance tests against the new stack, and update the Amazon Route 53 A record
C. Update the CloudFormation stack for the application tier using the T2 Unlimited option, run acceptance tests against the new stack, and update the Amazon Route 53 A record
D. Launch a new CloudFormation stack for all tiers of the application in a different region, run acceptance tests against the new stack, and update the Amazon Route 53 A record

Answer: C

NEW QUESTION 202
A DevOps Engineer is using AWS CodeDeploy across a fleet of Amazon EC2 instances in an EC2 Auto Scaling group. The associated CodeDeploy deployment group, which is integrated with EC2 Auto Scaling, is configured to perform in-place deployments with CodeDeployDefault.OneAtATime. During an ongoing new deployment, the Engineer discovers that, although the overall deployment finished successfully, two out of five instances have the previous application revision deployed. The other three instances have the newest application revision.
What is likely causing this issue?

A. The two affected instances failed to fetch the new deployment.
B. A failed AfterInstall lifecycle event hook caused the CodeDeploy agent to roll back to the previous version on the affected instances.
C. The CodeDeploy agent was not installed in two affected instances.
D. EC2 Auto Scaling launched two new instances while the new deployment had not yet finished, causing the previous version to be deployed on the affected instances.

Answer: D

NEW QUESTION 205
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 210
A company is running an application on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones.
After a recent application update, users are getting HTTP 502 Bad Gateway errors from the application URL. The DevOps Engineer cannot analyze the problem because Auto Scaling is terminating all EC2 instances shortly after launch for being unhealthy.
What steps will allow the DevOps Engineer access to one of the unhealthy instances to troubleshoot the deployed application?

A. Create an image from the terminated instance and create a new instance from that imag
B. The Application team can then log into the new instance.
C. As soon as a new instance is created by AutoScaling, put the instance into a Standby state as this will prevent the instance from being terminated.
D. Add a lifecycle hook to your Auto Scaling group to move instances in the Terminating state to the Terminating:Wait state.
E. Edit the Auto Scaling group to enable termination protection as this will protect unhealthy instances from being terminated.

Answer: B

Explanation:
https://aws.amazon.com/blogs/aws/auto-scaling-update-lifecycle-standby-detach/

NEW QUESTION 212
A Development team is working on a serverless application in AWS. To quickly identify and remediate potential production issues, the team decides to roll out changes to a small number of users as a test before the full release. The DevOps Engineer must develop a solution to minimize downtime and impact.
Which of the following solutions should be used to meet the requirements? (Select TWO.)

A. Create an Application Load Balancer with two target group
B. Set up the Application Load Balancer for Amazon API Gateway private integratio
C. Associate one target group to the current version and the other target group to the new versio
D. Configure API Gateway to route 10% of incoming traffic to the new versio
E. As the new version becomes stable, configure API Gateway to send all traffic to the new version and detach the old version from the load balancer.
F. Create an alias for an AWS Lambda function pointing to both the current and new version
G. Configure the alias to route 10% of incoming traffic to the new versio
H. As the new version is considered stable, update the alias to route all traffic to the new version.
I. Create a failover record set in AWS Route 53 pointing to the AWS Lambda endpoints for the old and new version
J. Configure Route 53 to route 10% of incoming traffic to the new versio
K. As the new version becomes stable, update the DNS record to route all traffic to the new version.
L. Create an ELB Network Load Balancer with two target group
M. Set up the Network Load Balancer for Amazon API Gateway private integration Associate one target group with the current version and the other target group with the new versio
N. Configure the load balancer to route 10% of incoming traffic to the new versio
O. As the new version becomes stable, detach the old version from the load balancer.
P. In Amazon API Gateway, create a canary release deployment by adding canary settings to the stage of a regular deploymen
Q. Configure API Gateway to route 10% of the incoming traffic to the canary releas
R. As the canary release is considered stable, promote it to a production release.

Answer: BE

Explanation:
https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html
https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-traffic-shifting/

NEW QUESTION 217
A company has developed a static website hosted on an Amazon S3 bucket. The website is deployed using AWS CloudFormation. The CloudFormation template defines an S3 bucket and a custom resource that copies content into the bucket from a source location.
The company has decided that it needs to move the website to a new location, so the existing CloudFormation stack must be deleted and re-created. However, CloudFormation reports that the stack could not be deleted cleanly.
What is the MOST likely cause and how can the DevOps Engineer mitigate this problem for this and future versions of the website?

A. Deletion has failed because the S3 bucket has an active website configuratio
B. Modify the CloudFormation template to remove the Website Configuration property from the S3 bucket resource.
C. Deletion has failed because the S3 bucket is not empt
D. Modify the custom resource's AWS Lambda function code to recursively empty the bucket when is Delet
E. RequestType
F. Deletion has failed because the custom resource does not define a deletion polic
G. Add a Deletion Policy property to the custom resource definition with a value of RemoveOnDeletion.
H. Deletion has failed because the S3 bucket is not empt
I. Modify the S3 bucket resource in the CloudFormation template to add a Deletion Policy property with a value of Empty.

Answer: D

NEW QUESTION 221
The Deployment team has grown substantially in recent months and so has the number of projects that use separate code repositories. The current process involves configuring AWS CodePipeline manually, and there have been service limit alerts for the count of Amazon S3 buckets.
Which pipeline option will reduce S3 bucket sprawl alerts?

A. Combine the multiple separate code repositories into a single one, and deploy using a global AWS CodePipeline that has logic for each project.
B. Create new pipelines by using the AWS API or AWS CLI, and configure them to use a single global S3 bucket with separate prefixes for each project.
C. Create a new pipeline in a different region for each project to bypass the service limits for S3 buckets in a single region.
D. Create a new pipeline and for S3 bucket for each project by using the AWS API or AWS CLI to bypass the service limits for S3 buckets in a single account

Answer: A

NEW QUESTION 223
A company is using Amazon EC2 for various workloads. Company policy requires that instances be managed centrally to standardize configurations. These configurations include standard logging, metrics, security assessments, and weekly patching.
How can the company meet these requirements? (Select THREE.)

A. Use AWS Config to ensure all EC2 instances are managed by Amazon Inspector.
B. Use AWS Config to ensure all EC2 instances are managed by AWS Systems Manager.
C. Use AWS Systems Manager to install and manage Amazon Inspector, Systems Manager Patch Manager, and the Amazon CloudWatch agent on all instances.
D. Use Amazon Inspector to install and manage AWS Systems Manager, Systems Manager Patch Manager, and the Amazon CloudWatch agent on all instances.
E. Use AWS Systems Manager maintenance windows with Systems Manager Run Command to schedule Systems Manager Patch Manager task
F. Use the Amazon CloudWatch agent to schedule Amazon Inspector assessment runs.
G. Use AWS Systems Manager maintenance windows with Systems Manager Run Command to schedule Systems Manager Patch Manager task
H. Use Amazon CloudWatch Events to schedule Amazon Inspector assessment runs.

Answer: BDE

NEW QUESTION 225
A DevOps engineer is tasked with migrating Docker containers used for a workload to AWS. The solution must allow for changes to be deployed into development and test environments automatically by updating each container and checking it into a container registry Once the containers are pushed, they must be deployed automatically
Which solution will meet these requirements?

A. Store container images in Amazon S3. Run the containers in AWS Elastic Beanstalk using a multicontainer Docker environmen
B. Configure Elastic Beanstalk to redeploy the containers if it detectsa new version in Amazon S3.
C. Store container images in AWS Artifact Use AWS CodePipeline to trigger a deployment if a new container version is create
D. Use AWS CodeDeploy to deploy new containers to Amazon EKS.
E. Store container images in Amazon ECR Use AWS CodePipeline to trigger a deployment if a new container version is created Use AWS CodeDeploy to deploy the image to AWS Fargate.
F. Store container images in Docker Hub Install Docker on an Amazon EC2 instance and use AWS CodePipeline and AWS CodeDeploy to deploy any new containers

Answer: C

NEW QUESTION 228
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 233
A company has an application that is using a MySQL -compatible Amazon Aurora Multi-AZ DB cluster as the database A cross-Region read replica has been created for disaster recovery purposes A DevOps engineer wants to automate the promotion of the replica so it becomes the primary database instance in the event of a failure
Which solution will accomplish this?

A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoints Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to trigger an AWS Lambda function that will promote the replica instance as the master.
B. Create an Aurora custom endpoint to point to the primary database instance Configure the application to use this endpoint Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.
C. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and pout the application to the newly promoted instance Create an Amazon CloudWatch alarm to trigger this Lambda function after the failure event occurs
D. Store the Aurora endpoint in AWS Systems Manager Parameter Store Create an Amazon EventBridge (Amazon CloudWatch Events) event that detects the database failure and runs an AWS Lambdafunction to promote the replica instance and update the endpoint URL stored in AWS Systems Manager Parameter Store Code the application to reload the endpoint from Parameter Store if a database connection fails.

Answer: A

NEW QUESTION 234
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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