Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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

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

Answer: D

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

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

Answer: A

NEW QUESTION 3
A Security team requires all Amazon EBS volumes that are attached to an Amazon EC2 instance to have AWS Key Management Service (AWS KMS) encryption enabled. If encryption is not enabled, the company's policy requires the EBS volume to be detached and deleted. A DevOps Engineer must automate the detection and deletion of unencrypted EBS volumes. Which method should the Engineer use to accomplish this with the LEAST operational effort?

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

Answer: B

NEW QUESTION 4
A company wants 10 use AWS development tools to replace Its current bash deployment scripts. The company currently deploys a LAMP application to a group of Amazon EC2 instances behind an Application Load Balancer (ALB). During the deployments, the company unit tests the committed application, stops and starts services, unregisters and re-registers instances with the load balancer, and updates Me permissions. The company wants to maintain the same deployment functionality through the shift to using AWS services.
Which solution will meet these requirements?

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

L. Unregister and re-register theinstances in the AWS CodeDeploy deployment group with the AL
M. Update the appspec.yml file to update file permissions without a custom script.

Answer: D

NEW QUESTION 5
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 6
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 7
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

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

Answer: C

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

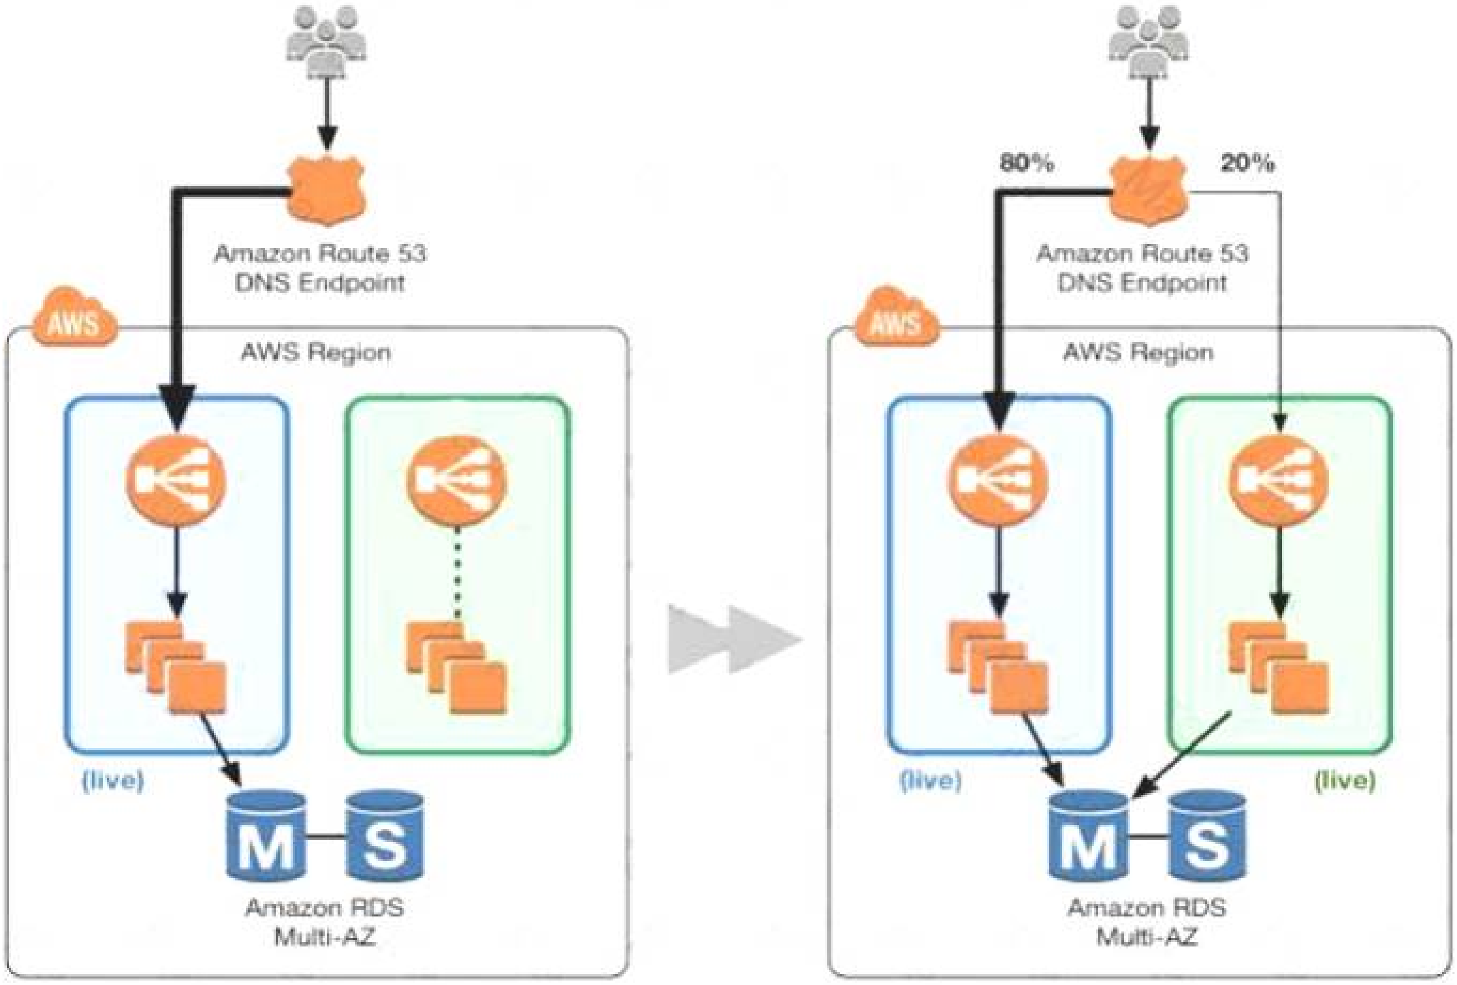

1) First create a new ELB which will be used to point to the new production changes.
2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80-20% traffic scenario. This is the normal case, the % can be changed based on the requirement.
3) Finally when all changes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

NEW QUESTION 8
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 9
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 DevOps engineer is troubleshooting deployments to a new application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. Instances sometimes come online before they are ready, which is leading to increased error rates among users. The current health check configuration gives instances a 60-second grace period and considers instances healthy after two 200 response codes from /index.php, a page that may respond intermittently during the deployment process. The development team wants instances to come online as soon as possible.
Which strategy would address this issue?

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

Answer: D

NEW QUESTION 10
A company recently launched an application that is more popular than expected. The company wants to ensure the application can scale to meet increasing demands and provide reliability using multiple Availability Zones (AZs) The application runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) A DevOps engineer has created an Auto Scaling group across multiple AZs for the application Instances launched in the newly added AZs are not receiving any traffic for the application.
What is likely causing this issue?

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

Answer: A

NEW QUESTION 12
A company is required to collect user consent to a privacy agreement. An application is deployed in six AWS Regions with two in North America, two in Europe, and two in Asia with a user base of 20-30 million users. The company needs to read and write data related to each user’s response, and ensure the responses are available in all six Regions.
What solution will satisfy these requirements while MINIMIZING latency?

A. Implement Amazon Aurora Global Database in each of the six Regions.
B. Implement Amazon DocumentDB (with MongoDB compatibility) in each of the six Regions.
C. Implement Amazon DynamoDB global tables in each of the six Regions.
D. Implement Amazon ElastiCache for Redis replication group in each of the six Regions.

Answer: C

NEW QUESTION 16
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 19
An Application team has three environments for their application: development, pre-production, and production. The team recently adopted AWS CodePipeline. However, the team has had several deployments of misconfigured or nonfunctional development code into the production environment, resulting in user disruption and downtime. The DevOps Engineer must review the pipeline and add steps to identify problems with the application before it is deployed.
What should the Engineer do to identify functional issues during the deployment process? (Choose two.)

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

Answer: BE

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

NEW QUESTION 23
A company wants to migrate a legacy application to AWS and develop a deployment pipeline that uses AWS services only. A DevOps engineer is migrating all of the application code from a Git repository to AWS CodeCommit while preserving the history of the repository. The DevOps engineer has set all the permissions within CodeCommit, installed the Git client and the AWS CLI on a local computer, and is ready to migrate the repository.
Which actions will follow?

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

Answer: D

NEW QUESTION 25
A development team manually builds an artifact locally and then places it in an Amazon S3 bucket. The application has a local cache that must be cleared when a deployment occurs. The team executes a command to do this, downloads the artifact from Amazon S3, and unzips the artifact to complete the deployment.
A DevOps team wants to migrate to a CI/CD process and build in checks to stop and roll back the deployment when a failure occurs. This requires the team to track the progression ot the deployment.
Which combination of actions will accomplish this? (Select THREE.)

A. Allow developers to check the code into a code repositor
B. Using Amazon CloudWatch Events, on every pull into master, trigger an AWS Lambda function to build the artifact and store it in Amazon
C. Create a custom script to clear the cach
D. Specify the script in the BeforeInstall lifecycle hook in the AppSpec file.
E. Create user data for each Amazon EC2 instance that contains the clear cache scrip
F. Once deployed, test the applicatio
G. If it is not successful, deploy it again.
H. Set up AWS CodePipeline to deploy the applicatio
I. Allow developers to check the code into a code repository as a source for the pipeline.
J. Use AWS CodeBuild to build the artifact and place it in Amazon S3. Use AWS CodeDeploy to deploy the artifact to Amazon EC2 instances.
K. Use AWS Systems Manager to fetch the artifact from Amazon S3 and deploy it to all the instances.

Answer: ADE

NEW QUESTION 27
A company is testing a web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company uses a blue/green deployment process with immutable instances when deploying new software.
During testing, users are being automatically logged out of the application at random times. Testers also report that, when a new version of the application is deployed, all users are logged out. The Development team needs a solution to ensure users remain logged in across scaling events and application deployments.
What is the MOST efficient way to ensure users remain logged in?

A. Enable smart sessions on the load balancer and modify the application to check for an existing session.
B. Enable session sharing on the load balancer and modify the application to read from the session store.
C. Store user session information in an Amazon S3 bucket and modify the application to read session information from the bucket.
D. Modify the application to store user session information in an Amazon ElastiCache cluser.

Answer: D

NEW QUESTION 28
A Security team is concerned that a Developer can unintentionally attach an Elastic IP address to an Amazon EC2 instance in production. No Developer should be allowed to attach an Elastic IP address to an instance. The Security team must be notified if any production server has an Elastic IP address at any time.
How can this task be automated?

A. Use Amazon Athena to query AWS CloudTrail logs to check for any associate-address attempt
B. Create an AWS Lambda function to dissociate the Elastic IP address from the instance, and alert the Security team.
C. Attach an IAM policy to the Developer's IAM group to deny associate-address permission
D. Create a custom AWS Config rule to check whether an Elastic IP address is associated with any instance tagged as production, and alert the Security team.
E. Ensure that all IAM groups are associated with Developers do not have associate-address permissions.Create a scheduled AWS Lambda function to check whether an Elastic IP address is associated with any instance tagged as production, and alert the Security team if an instance has an Elastic IP address associated with it.
F. Create an AWS Config rule to check that all production instances have the EC2 IAM roles that include deny associate-address permission
G. Verify whether there is an Elastic IP address associated with any instance, and alert the Security team if an instance has an Elastic IP address associated with it.

Answer: B

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-migrate-ipv6.html#vpc-migrate-ipv6-sg-rules

NEW QUESTION 29
A company is deploying a new application that uses Amazon EC2 instances. The company needs a solution to query application logs and AWS account API activity. Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 33
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 38
A company has built a web service that runs on Amazon EC2 instances behind an Application Load Balancer (ALB) the company has deployed the application in us-east-1 Amazon Route 53 provides an external DNS that routes traffic from example.com to the application, created with appropriate health checks.
The company has deployed a second environment for the application in eu-west-1 the company wants traffic to be routed to whichever environment results m the best response time for each user. If there is an outage in one Region, traffic should be directed to the other environment.
Which configuration will achieve this requirements?

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

Answer: B

NEW QUESTION 41
A legacy web application stores access logs in a proprietary text format. One of the security requirements is to search application access events and correlate them with access data from many different systems. These searches should be near-real time.
Which solution offloads the processing load on the application server and provides a mechanism to search the data in near-real time?

A. Install the Amazon CloudWatch Logs agent on the application server and use CloudWatch Events rules to search logs for access event
B. Use Amazon CloudSearch as an interface to search for events.
C. Use the third-party file-input plugin Logstash to monitor the application log file, then use a custom dissect filter on the agent to parse the log entries into the JSON forma
D. Output the events to Amazon ES to be searche
E. Use the Elasticsearch API for querying the data.
F. Upload the log files to Amazon S3 by using the S3 sync comman
G. Use Amazon Athena to define the structure of the data as a table, with Athena SQL queries to search for access events.
H. Install the Amazon Kinesis Agent on the application server, configure it to monitor the log files, and send it to a Kinesis strea
I. Configure Kinesis to transform the data by using an AWS Lambda function, and forward events to Amazon ES for analysi
J. Use the Elasticsearch API for querying the data.

Answer: D

Explanation:
https://docs.aws.amazon.com/zh_cn/streams/latest/dev/writing-with-agents.html

NEW QUESTION 46
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 48
A company is implementing an Amazon ECS cluster to run its workload. The company architecture will run multiple ECS services on the cluster, with an Application Load Balancer on the front end, using multiple target groups to route traffic. The Application Development team has been struggling to collect logs that must be collected and sent to an Amazon S3 bucket for near-real time analysis
What must the DevOps Engineer configure in the deployment to meet these requirements? (Select THREE)

A. Install the Amazon CloudWatch Logs logging agent on the ECS instance
B. Change the logging driver in the ECS task definition to 'awslogs'.
C. Download the Amazon CloudWatch Logs container instance from AWS and configure it as a task.Update the application service definitions to include the logging task.
D. Use Amazon CloudWatch Events to schedule an AWS Lambda function that will run every 60 seconds running the create-export -task CloudWatch Logs command, then point the output to the logging S3 bucket.
E. Enable access logging on the Application Load Balancer, then point it directly to the S3 logging bucket.
F. Enable access logging on the target groups that are used by the ECS services, then point it directly to the S3 logging bucket.
G. Create an Amazon Kinesis Data Firehose with a destination of the S3 logging bucket, then create an Amazon CloudWatch Logs subscription filter for Kinesis

Answer: BDF

NEW QUESTION 52
A company must ensure consistent behavior of an application running on Amazon Linux in its corporate
ecosystem before moving into AWS. The company has an existing automated server build system using VMware. The goal is to demonstrate the functionality of the application and its prerequisites on the new target operating system.
The DevOps Engineer needs to use the existing corporate server pipeline and virtualization software to create a server image. The server image will be tested on-premises to resemble the build on Amazon EC2 as closely as possible.
How can this be accomplished?

A. Download and integrate the latest ISO of CentOS 7 and execute the application deployment on the resulting server.
B. Launch an Amazon Linux AMI using an AWS OpsWorks deployment agent onto the on-premises infrastructure, then execute the application deployment.
C. Build an EC2 instance with the latest Amazon Linux operating system, and use the AWS Import/Export service to export the EC2 image to a VMware ISO in Amazon S3. Then import the resulting ISO onto the on-premises system.
D. Download and integrate the latest ISO of Amazon Linux 2 and execute the application deployment on the resulting serve
E. Confirm that operating system testing results are consistent with EC2 operating system behavior.

Answer: D

NEW QUESTION 56
A government agency has multiple AWS accounts, many of which store sensitive citizen information. A Security team wants to detect anomalous account and

network activities (such as SSH brute force attacks) in any account and centralize that information in a dedicated security account. Event information should be stored in an Amazon S3 bucket in the security account, which is monitored by the department's Security Information and Even Manager (SIEM) system.
How can this be accomplished?

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

Answer: C

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

NEW QUESTION 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 61
A Development team is building more than 40 applications. Each app is a three-tiered web application based on an ELB Application Load Balancer, Amazon EC2, and Amazon RDS. Because the applications will be used internally, the Security team wants to allow access to the 40 applications only from the corporate network and block access from external IP addresses. The corporate network reaches the internet through proxy servers. The proxy servers have 12 proxy IP addresses that are being changed one or two times per month. The Network Infrastructure team manages the proxy servers; they upload the file that contains the latest proxy IP addresses into an Amazon S3 bucket. The DevOps Engineer must build a solution to ensure that the applications are accessible from the corporate network.
Which solution achieves these requirements with MINIMAL impact to application development, MINIMAL operational effort, and the LOWEST infrastructure cost?

A. Implement an AWS Lambda function to read the list of proxy IP addresses from the S3 object and to update the ELB security groups to allow HTTPS only from the given IP addresse
B. Configure the S3 bucket to invoke the Lambda function when the object is update
C. Save the IP address list to the S3 bucket when they are changed.
D. Ensure that all the applications are hosted in the same Virtual Private Cloud (VPC). Otherwise, consolidate the applications into a single VP
E. Establish an AWS Direct Connect connection with an active/standby configuratio
F. Change the ELB security groups to allow only inbound HTTPS connections from the corporate network IP addresses.
G. Implement a Python script with the AWS SDK for Python (Boto), which downloads the S3 object that contains the proxy IP addresses, scans the ELB security groups, and updates them to allow only HTTPS inbound from the given IP addresse
H. Launch an EC2 instance and store the script in the instanc
I. Use a cron job to execute the script daily.
J. Enable ELB security groups to allow HTTPS inbound access from the Interne
K. Use Amazon Cognito to integrate the company's Active Directory as the identity provide
L. Change the 40 applications to integrate with Amazon Cognito so that only company employees can log into the applicatio
M. Save the user access logs to Amazon CloudWatch Logs to record user access activities

Answer: A

NEW QUESTION 63
A company is using AWS CodeCommit as its source code repository. After an internal audit, the compliance team mandates that any code change that go into the master branch must be committed by senior developers.
Which solution will meet these requirements?

A. Create two repositories in CodeCommit: one for working and another for the maste
B. Create separate IAM groups for senior developers and developer
C. Assign the resource-level permissions on the repositories tied to the IAM group
D. After the code changes are reviewed, sync the approved files to the master code commit repository.
E. Create a repository in CodeCommi
F. Create separate IAM groups for senior developers and developers.Assign code commit permissions for both groups, with code merge permissions for the senior developers grou
G. Create a trigger to notify senior developers with a URL link to approve or deny commit requests delivered through Amazon SN
H. Once a senior developer approves the code, the code gets merged to the master branch.
I. Create a repository in CodeCommit with a working and master branc
J. Create separate IAM groups for senior developers and developer

K. Use an IAM policy to assign each IAM group their corresponding branche
L. Once the code is merged to the working branch, senior developers can pull the changes from the working branch to the master branch.
M. Create a repository in CodeCommi
N. Create separate IAM groups for senior developers and developers.Use AWS Lambda triggers on the master branch and get the user name of the developer at the event object of the Lambda functio
O. Validate the user name with the IAM group to approve or deny the commit.

Answer: C

NEW QUESTION 66
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 70
A company runs an application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones in us-east1. The application stores data in an Amazon RDS MySQL Multi-AZ DB instance.
A DevOps Engineer wants to modify the current solution and create a hot standby of the environment in another region to minimize downtime if a problem occurs in us-east-1.
Which combination of steps should the DevOps Engineer take to meet these requirements? (Select THREE.)

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

Answer: ABE

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

NEW QUESTION 73
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 74
A DevOps Engineer is designing a deployment strategy for a web application. The application will use an Auto Scaling group to launch Amazon EC2 instances using an AMI. The same infrastructure will be deployed in multiple environments (development, test, and quality assurance). The deployment strategy should meet the following requirements: "¢ Minimize the startup time for the instance "¢ Allow the same AMI to work in multiple environments "¢ Store secrets for multiple environments securely
How should this be accomplished?

A. Preconfigure the AMI using an AWS Lambda function that launches an Amazon EC2 instance, and then runs a script to install the software and create the AM
B. Configure an Auto Scaling lifecycle hook to determine which environment the instance is launched in, and, based on that finding, run a configuration scrip
C. Save the secrets on an .ini file and store them in Amazon S3. Retrieve the secrets using a configuration script in EC2 user data.
D. Preconfigure the AMI by installing all the software using AWS Systems Manager automation and configure Auto Scaling to tag the instances at launch with their specific environmen
E. Then use a bootstrap script in user data to read the tags and configure settings for the environmen
F. Use the AWS Systems Manager Parameter Store to store the secrets using AWS KMS.
G. Use a standard AMI from the AWS Marketplac
H. Configure Auto Scaling to detect the current environmen

I. Install the software using a script in Amazon EC2 user dat
J. Use AWS Secrets Manager to store the credentials for all environments.
K. Preconfigure the AMI by installing all the software and configuration for all environment
L. Configure Auto Scaling to tag the instances at launch with their environmen
M. Use the Amazon EC2 user data to trigger an AWS Lambda function that reads the instance ID and then reconfigures the setting for the proper environmen
N. Use the AWS Systems Manager Parameter Store to store the secrets using AWS KMS.

Answer: A

NEW QUESTION 76
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 79
A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project.
How can this issue be corrected in the MOST secure manner?

A. Add the bucket name to the AllowedBuckets section of the CodeBuild project setting
B. Update the build spec to use the AWS CLI to download the databasepopulation script.
C. Modify the S3 bucket settings to enable HTTPS basic authentication and specify a toke
D. Update the build spec to use cURL to pass the token and download the database population script.
E. Remove unauthenticated access from the S3 bucket with a bucket polic
F. Modify the service role for the CodeBuild project to include Amazon S3 acces
G. Use the AWS CLI to download the database population script.
H. Remove unauthenticated access from the S3 bucket with a bucket polic
I. Use the AWS CLI to download the database population script using an IAM access key and a secret access key.

Answer: C

NEW QUESTION 83
A DevOps engineer is creating a CI/CD pipeline for an Amazon ECS service. The ECS container instances run behind an Application Load Balancer as the web tier of a three-tier application. An acceptance criterion (or a successful deployment is the verification that the web tier can communicate with the database and middleware tiers of the application upon deployment.
How can this be accomplished in an automated fashion?

A. Create a health check endpoint in the web application that tests connectivity to the data and middleware tier
B. Use this endpoint as the health check URL for the load balancer.
C. Create an approval step for the quality assurance team to validate connectivit
D. Reject changes in the pipeline if there is an issue with connecting to the dependent tiers.
E. Use an Amazon RDS active connection count and an Amazon CloudWatch ELB metric to alarm on a significant change to the number of open connections.
F. Use Amazon Route 53 health checks to detect issues with the web service and roll back the CI/CD pipeline if there is an error.

Answer: A

NEW QUESTION 85

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

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

Answer: C

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

NEW QUESTION 90
A company's web application will be migrated to AWS. The application is designed so that there is no server-side code required. As part of the migration, the company would like to improve the security of the application by adding HTTP response headers, following the Open Web Application Security Project (OWASP) secure headers recommendations.
How can this solution be implemented to meet the security requirements using best practices?

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

Answer: C

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

NEW QUESTION 95
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 102
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 105
A company has multiple development teams sharing one AWS account. The development team's manager wants to be able to automatically stop Amazon EC2 instances and receive notifications if resources are idle and not tagged as production resources
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 110
A company's security team discovers that IAM access keys were exposed in a public code repository. Moving forward, the DevOps team wants to implement a solution that will automatically disable any keys that are suspected of being compromised, and notify the security team.
Which solution will accomplish this?

A. Create an Amazon CloudWatch Events event for Amazon Mad
B. Create an Amazon SNS topic with two subscriptions: one to notify the security team and another to trigger an AWS Lambda function that disables the access keys.
C. Enable Amazon GuardDuty and set up an Amazon CloudWatch Events rule event for GuardDuty.Trigger an AWS Lambda function to check if the event relates to compromised key
D. If so, send a notification to the security team and disable the access keys.
E. Run an AWS CloudWatch Events rule every 5 minutes to invoke an AWS Lambda function that checks to see if the compromised tag for any access key is set to tru
F. If s
G. notify the security team and disable the access keys.
H. Set up AWS Config and create an AWS CloudTrail event for AWS Confi
I. Create an Amazon SNS topic with two subscriptions: one to notify the security team and another to trigger an AWS Lambda function that disables the access keys.

Answer: B

NEW QUESTION 112
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
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 116
A company's application is running on Amazon EC2 instances in an Auto Scaling group. A DevOps engineer needs to ensure there are at least four application servers running at all times. Whenever an update has to be made to the application, the engineer creates a new AMI with the updated configuration and updates the AWS CloudFormation template with the new AMI ID. After the stack update finishes, the engineer manually terminates the old instances one by one. verifying that the new instance is operational before proceeding. The engineer needs to automate this process.
Which action will allow for the LEAST number of manual steps moving forward?

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

Answer: A

NEW QUESTION 121
A company is using AWS CodeBuild, AWS CodeDeploy, and AWS CodePipeline to deploy applications automatically to an Amazon EC2 instance. A DevOps Engineer needs to perform a security assessment scan of the operating system on every application deployment to the environment.
How should this be automated?

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

Answer: A

NEW QUESTION 122
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 127
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 130
An IT department manages a portfolio with Windows and Linux (Amazon and Red Hat Enterprise Linux) servers both on-premises and on AWS. An audit reveals that there is no process for updating OS and core application patches, and that the servers have inconsistent patch levels.
Which of the following provides the MOST reliable and consistent mechanism for updating and maintaining all servers at the recent OS and core application patch levels?

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

Answer: A

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

NEW QUESTION 131
A DevOps engineer must ensure all IAM entity configurations across multiple AWS accounts in AWS Organizations are compliant with corporate IAM policies.
Which combination of steps will accomplish this? (Select TWO.)

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

Answer: BE

NEW QUESTION 136
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 137
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 140
An education company has a Docker-based application running on multiple Amazon EC2 instances in an Amazon ECS cluster. When deploying a new version of the application, the Developer, pushes a new image to a private Docker container registry, and then stops and starts all tasks to ensure that they all have the latest version of the application. The Developer discovers that the new tasks are occasionally running with an old image.
How can this issue be prevented?

A. After pushing the new image, restart ECS Agent, and then start the tasks.
B. Use "latest" for the Docker image tag in the task definition.
C. Update the digest on the task definition when pushing the new image.
D. Use Amazon ECR for a Docker container registry.

Answer: C

NEW QUESTION 144
A company wants to adopt a methodology for handling security threats from leaked and compromised IAM access keys. The DevOps Engineer has been asked to automate the process of acting upon compromised access keys, which includes identifying users, revoking their permissions, and sending a notification to the Security team.
Which of the following would achieve this goal?

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

Answer: B

NEW QUESTION 147
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 152
A company mandates the creation of capture logs for everything running in its AWS account. The account has multiple VPCs with Amazon EC2 instances, Application Load Balancers, Amazon RDS MySQL databases, and AWS WAF rules configured. The logs must be protected from deletion. A daily visual analysis of log anomalies from the previous day is required.
Which combination of actions should a DevOps Engineer take to accomplish this? (Choose three.)

A. Configure an AWS Lambda function to send all CloudWatch logs to an Amazon S3 bucke
B. Create a dashboard report in Amazon QuickSight.
C. Configure AWS CloudTrail to send all logs to Amazon Inspecto
D. Create a dashboard report in Amazon QuickSight.
E. Configure Amazon S3 MFA Delete on the logging Amazon S3 bucket.
F. Configure an Amazon S3 object lock legal hold on the logging Amazon S3 bucket.
G. Configure AWS Artifact to send all logs to the logging Amazon S3 bucke
H. Create a dashboard report in Amazon QuickSight.
I. Deploy an Amazon CloudWatch agent to all Amazon EC2 instances.

Answer: ADF

NEW QUESTION 153
The management team at a company with a large on-premises OpenStack environment wants to move non-production workloads to AWS. An AWS Direct Connect connection has been provisioned and configured to connect the environments. Due to contractual obligations, the production workloads must remain on-premises, and will be moved to AWS after the next contract negotiation. The company follows Center for Internet Security (CIS) standards for hardening images; this configuration was developed using the company's configuration management system.
Which solution will automatically create an identical image in the AWS environment without significant overhead?

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

Answer: D

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

NEW QUESTION 156
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 160
A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real time and 1% of requests should route to the secondary region to continuously verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the secondary region must be able to scale up to handle all traffic.

How should a DevOps Engineer meet these requirements?

A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session dat
B. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.
C. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session dat
D. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.
E. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS PostgreSQL with cross-region replication for session dat
F. Deploy the web applicationwith client-side logic to call the API Gateway directly.
G. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session dat
H. Enable an Amazon CloudFront weighted distribution across region
I. Point the Amazon Route 53 DNS record at the CloudFront distribution.

Answer: B

NEW QUESTION 162
A company wants to use a grid system for a proprietary enterprise in-memory data store on top of AWS. This system can run in multiple server nodes in any Linux-based distribution. The system must be able to reconfigure the entire cluster every time a node is added or removed. When adding or removing nodes, an / etc./cluster/nodes.config file must be updated, listing the IP addresses of the current node members of that cluster
The company wants to automate the task of adding new nodes to a cluster. What can a DevOps Engineer do to meet these requirements?

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

Answer: A

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

NEW QUESTION 166
A DevOps Engineer manages an application that has a cross-region failover requirement. The application stores its data in an Amazon Aurora on Amazon RDS database in the primary region with a read replica in the secondary region. The application uses Amazon Route 53 to direct customer traffic to the active region.
Which steps should be taken to MINIMIZE downtime if a primary database fails?

A. Use Amazon CloudWatch to monitor the status of the RDS instanc
B. In the event of a failure, use a CloudWatch Events rule to send a short message service (SMS) to the Systems Operator using Amazon SN
C. Have the Systems Operator redirect traffic to an Amazon S3 static website that displays a downtime messag
D. Promote the RDS read replica to the maste
E. Confirm that the application is working normally, then redirect traffic from the Amazon S3 website to the secondary region.
F. Use RDS Event Notification to publish status updates to an Amazon SNS topi
G. Use an AWS Lambda function subscribed to the topic to monitor database healt
H. In the event of a failure, the Lambda function promotes the read replica, then updates Route 53 to redirect traffic from the primary region to the secondary region.
I. Set up an Amazon CloudWatch Events rule to periodically invoke an AWS Lambda function that checks the health of the primary databas
J. If a failure is detected, the Lambda function promotes the read replic
K. Then, update Route 53 to redirect traffic from the primary to the secondary region.
L. Set up Route 53 to balance traffic between both regions equall
M. Enable the Aurora multi-master option, then set up a Route 53 health check to analyze the health of the database
N. Configure Route 53 to automatically direct all traffic to the secondary region when a primary database fails.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.html

NEW QUESTION 167
A DevOps Engineer manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The Engineer needs to implement a deployment strategy that:
Launches a second fleet of instances with the same capacity as the original fleet. Maintains the original fleet unchanged while the second fleet is launched. Transitions traffic to the second fleet when the second fleet is fully deployed. Terminates the original fleet automatically 1 hour after transition.
Which solution will satisfy these requirements?

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

G. Use AWS Elastic Beanstalk with the configuration set to Immutabl
H. Create an .ebextension using the Resources key that sets the deletion policy of the ALB to 1 hour, and deploy the application.

Answer: B

NEW QUESTION 172
A web application for healthcare services runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. A DevOps Engineer must create a mechanism in which an EC2 instance can be taken out of production so its system logs can be analyzed for issues to quickly troubleshot problems on the web tier.
How can the Engineer accomplish this task while ensuring availability and minimizing downtime?

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

Answer: D

NEW QUESTION 176
A company has 100 GB of log data in an Amazon S3 bucket stored in .csv format. SQL developers want to query this data and generate graphs to visualize it. They also need an efficient, automated way to store metadata from the .csv file.
Which combination of steps should be taken to meet these requirements with the LEAST amount of effort? (Select THREE.)

A. Filter the data through AWS X-Ray to visualize the data.
B. Filter the data through Amazon QuickSight to visualize the data.
C. Query the data with Amazon Athena.
D. Query the data with Amazon Redshift.
E. Use AWS Glue as the persistent metadata store.
F. Use Amazon S3 as the persistent metadata store.

Answer: BCE

NEW QUESTION 180
A company updated the AWS CloudFormation template tor a critical business application. The stack update process Tailed due to an error in me updated template, and CloudFormation automatically began the stack rollback process Later, a DevOps engineer found the application was still unavailable, and that the stack was in the UPDATE_ROLLBACK_FALED state
Which combination of actions will allow the stack rollback to complete successful/? (Select TWO)

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

Answer: AB

NEW QUESTION 185
A company wants to use AWS CloudFormation for infrastructure deployment. The company has strict tagging and resource requirements and wants to limit the deployment to two Regions. Developers will need to deploy multiple versions of the same application.
Which solution ensures resources are deployed in accordance with company policy?

A. Create AWS Trusted Advisor checks to find and remediate unapproved CloudFormation StackSets.
B. Create a CloudFormation drift detection operation to find and remediate unapproved CloudFormation StackSets.
C. Create CloudFormation StackSets with approved CloudFormation templates.
D. Create AWS Service Catalog products with approved CloudFormation templates.

Answer: C

NEW QUESTION 190
A large enterprise is deploying a web application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS Oracle DB instance and Amazon DynamoDB. There are separate environments for development, testing, and production.
What is the MOST secure and flexible way to obtain password credentials during deployment?

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

Answer: B

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

NEW QUESTION 195
You have decided that you need to change the instance type of your production instances which are running as part of an AutoScaling group. The entire architecture is deployed using CloudFormation Template. You currently have 4 instances in Production. You cannot have any interruption in service and need to ensure 2 instances are always runningduring the update? Which of the options below listed can be used for this?

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

Answer: A

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

- https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/

NEW QUESTION 198
A DevOps engineer is currently running a container-based workload on-premises The engineer wants to move the application to AWS, but needs to keep the on-premises solution active because not all APIs will move at the same time. The traffic between AWS and the on-premises network should be secure and encrypted at all times. Low management overload is also a requirement.
Which combination of actions will meet these criteria? (Select THREE.)

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

Answer: BDF

NEW QUESTION 200
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 202
A company uses a complex system that consists of networking, IAM policies, and multiple three-tier applications. Requirements are still being defined for a new system, so the number of AWS components present in the final design is not known. The DevOps Engineer needs to begin defining AWS resources using AWS CloudFormation to automate and version-control the new infrastructure.
What is the best practice for using CloudFormation to create new environments?

A. Manually construct the networking layer using Amazon VPC and then define all other resources using CloudFormation.
B. Create a single template to encompass all resources that are required for the system so there is only one template to version-control.
C. Create multiple separate templates for each logical part of the system, use cross-stack references in CloudFormation, and maintain several templates in version control.
D. Create many separate templates for each logical part of the system, and provide the outputs from one to the next using an Amazon EC2 instance running SDK for granular control.

Answer: C

NEW QUESTION 205
A DevOps Engineer at a startup cloud-based gaming company has the task formalizing deployment strategies. The strategies must meet the following requirements:
Use standard Git commands, such as git clone and git push for the code repository. Management tools should maximize the use of platform solutions where possible. Deployment packages must be immutable and in the form of Docker images.
How can the Engineer meet these requirements?

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

Answer: B

NEW QUESTION 208
A Development team uses AWS CodeCommit for source code control. Developers apply their changes to various feature branches and create pull requests to move those changes to the master branch when they are ready for production. A direct push to the master branch should not be allowed. The team applied the AWS managed policy AWSCodeCommitPowerUser to the Developers' IAM Rote, but now members are able to push to the master branch directly on every repository in the AWS account.
What actions should be taken to restrict this?

A. Create an additional policy to include a deny rule for the codecommit:GitPush action, and include arestriction for the specific repositories in the resource statement with a condition for the master reference.
B. Remove the IAM policy and add an AWSCodeCommitReadOnly polic
C. Add an allow rule for the codecommit:GitPush action for the specific repositories in the resource statement with a condition for the master reference.
D. Modify the IAM policy and include a deny rule for the codecommit:GitPush action for the specific repositories in the resource statement with a condition for the master reference.
E. Create an additional policy to include an allow rule for the codecommit:GitPush action and include a restriction for the specific repositories in the resource statement with a condition for the feature branches reference.

Answer: A

Explanation:
https://aws.amazon.com/pt/blogs/devops/refining-access-to-branches-in-aws-codecommit/

NEW QUESTION 210
An Information Security policy requires that all publicly accessible systems be patched with critical OS security patches within 24 hours of a patch release. All instances are tagged with the Patch Group key set to 0. Two new AWS Systems Manager patch baselines for Windows and Red Hat Enterprise Linux (RHEL) with zero-day delay for security patches of critical severity were created with an auto-approval rule. Patch Group 0 has been associated with the new patch baselines.
Which two steps will automate patch compliance and reporting? (Select TWO.)

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

Answer: AC

NEW QUESTION 212
A DevOps Engineer is working on a project that is hosted on Amazon Linux and has failed a security review. The DevOps Manager has been asked to review the company buildspec.yaml file for an AWS CodeBuild project and provide recommendations. The buildspec.yaml file is configured as follows:

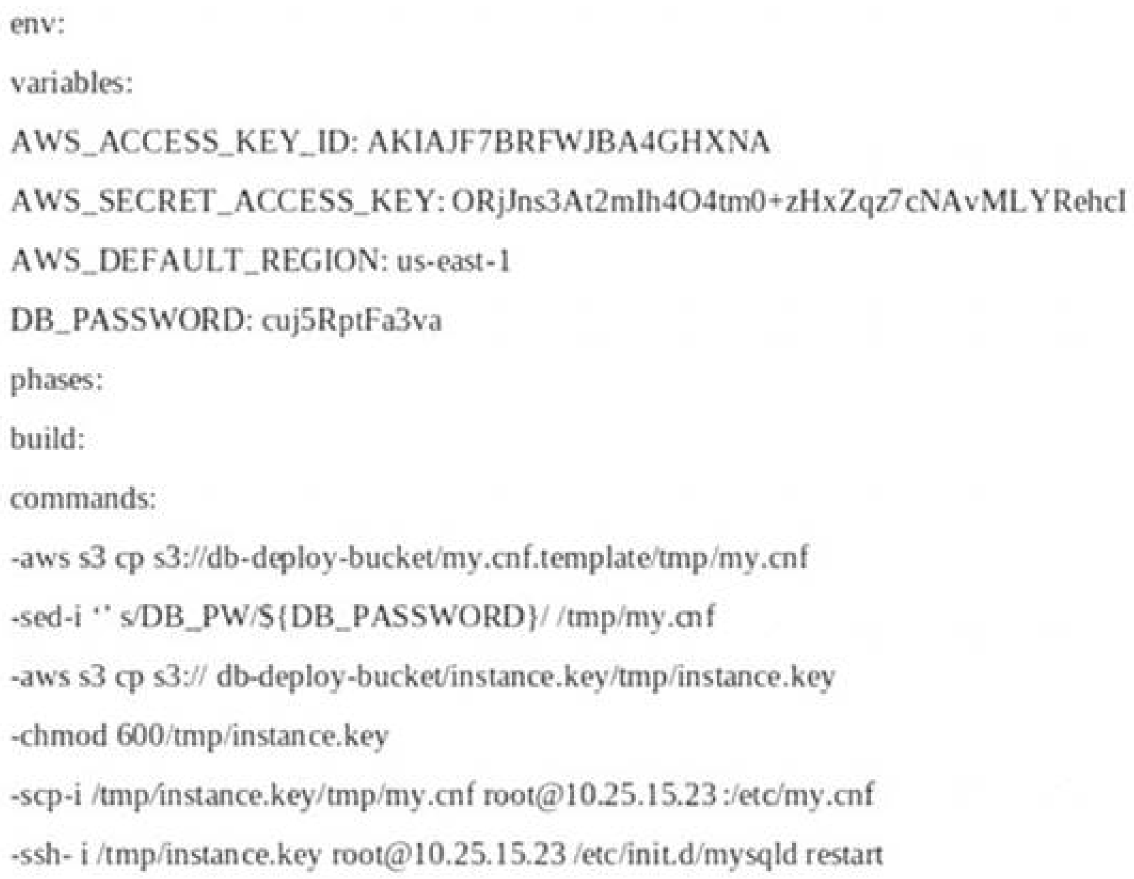

```
env:
variables:
AWS_ACCESS_KEY_ID: AKIAJF7BRFWJBA4GHXNA
AWS_SECRET_ACCESS_KEY: ORjJns3At2mIh4O4tm0+zHxZqz7cNAvMLYRehcl
AWS_DEFAULT_REGION: us-east-1
DB_PASSWORD: cuj5RptFa3va
phases:
build:
commands:
-aws s3 cp s3://db-deploy-bucket/my.cnf.template/tmp/my.cnf
-sed-i ‘’ s/DB_PW/${DB_PASSWORD}/ /tmp/my.cnf
-aws s3 cp s3:// db-deploy-bucket/instance.key/tmp/instance.key
-chmod 600/tmp/instance.key
-scp-i /tmp/instance.key/tmp/my.cnf root@10.25.15.23 :/etc/my.cnf
-ssh- i /tmp/instance.key root@10.25.15.23 /etc/init.d/mysqld restart
```

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

A. Add a post-build command to remove the temporary files from the container before termination to ensure they cannot be seen by other CodeBuild users.
B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable.
C. Store the DB_PASSWORD as a SecureString value in AWS Systems Manager Parameter Store and then remove the DB_PASSWORD from the environment variables.
D. Move the environment variables to the '˜db-deploy-bucket' Amazon S3 bucket, add a prebuild stage to download, then export the variables.
E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.
F. Scramble the environment variables using XOR followed by Base64, add a section to install, and then run XOR and Base64 to the build phase.

Answer: BCE

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

NEW QUESTION 216
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 218
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 219
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 223
A company is using AWS for an application. The Development team must automate its deployments. The team has set up an AWS CodePipeline to deploy the application to Amazon EC2 instances by using AWS CodeDeploy after it has been built using the AWS CodeBuild service.
The team would like to add automated testing to the pipeline to confirm that the application is healthy before deploying it to the next stage of the pipeline using the same code. The team requires a manual approval action before the application is deployed, even if the test is successful. The testing and approval must be accomplished at the lowest costs, using the simplest management solution.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 228
A DevOps Engineer uses Docker container technology to build an image-analysis application. The application often sees spikes in traffic. The Engineer must automatically scale the application in response to customer demand while maintaining cost effectiveness and minimizing any impact on availability.
What will allow the FASTEST response to spikes in traffic while fulfilling the other requirements?

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

Answer: A

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

NEW QUESTION 231
A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps Engineer is tasked with minimizing application response times and improving availability for users in both Regions.
Which combination of actions should be taken to address the latency issues? (Choose three.)

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

Answer: CDF

NEW QUESTION 236
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 239

......

Relate Links

100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with Exambible Prep Materials

https://www.exambible.com/AWS-Certified-DevOps-Engineer-Professional-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/