

# Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

<https://www.2passeasy.com/dumps/AWS-Certified-DevOps-Engineer-Professional/>



#### NEW QUESTION 1

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

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

**Answer:** C

#### NEW QUESTION 2

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 3

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 master.
- B. Allow Developers to use feature branches and create a pull request when a feature is complete.
- 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 complete.
- 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 master.
- 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 complete.
- 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 master.
- 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 master.
- K. Allow Developers to use feature branches and create a pull request when a feature is complete.
- L. Allow an approver to use CodeCommit to view the changes and approve the pull requests.

**Answer:** D

#### NEW QUESTION 4

A company is setting up a centralized logging solution on AWS and has several requirements. The company wants its Amazon CloudWatch Logs and VPC Flow logs to come from different sub accounts and to be delivered to a single auditing account. However, the number of sub accounts keeps changing. The company also needs to index the logs in the auditing account to gather actionable insight. How should a DevOps Engineer implement the solution to meet all of the company's requirements?

- A. Use AWS Lambda to write logs to Amazon ES in the auditing account.
- B. Create an Amazon CloudWatch subscription filter and use Amazon Kinesis Data Streams in the sub accounts to stream the logs to the Lambda function deployed in the auditing account.
- C. Use Amazon Kinesis Streams to write logs to Amazon ES in the auditing account.
- D. Create a CloudWatch subscription filter and use Kinesis Data Streams in the sub accounts to stream the logs to the Kinesis stream in the auditing account.
- E. Use Amazon Kinesis Firehose with Kinesis Data Streams to write logs to Amazon ES in the auditing account.
- F. Create a CloudWatch subscription filter and stream logs from sub accounts to the Kinesis stream in the auditing account.
- G. Use AWS Lambda to write logs to Amazon ES in the auditing account.
- H. Create a CloudWatch subscription filter and use Lambda in the sub accounts to stream the logs to the Lambda function deployed in the auditing account.

**Answer:** C

#### Explanation:

<https://aws.amazon.com/pt/blogs/architecture/central-logging-in-multi-account-environments/>

#### 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 launched
- 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 environment
- 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 highly regulated company has a policy that DevOps Engineers should not log in to their Amazon EC2 instances except in emergencies. If a DevOps Engineer does log in, the Security team must be notified within 15 minutes of the occurrence.

Which solution will meet these requirements?

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

**Answer:** B

#### NEW QUESTION 7

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

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

**Answer:** C

#### Explanation:

The ideal way is to create a new launch configuration, attach it to the existing Auto Scaling group, and terminate the running instances.

Option A is invalid because Elastic beanstalk cannot launch new instances on demand. Since the current scenario requires Autoscaling, this is not the ideal option

Option B is invalid because this will be a maintenance overhead, since you just have an Autoscaling Group.

There is no need to create a whole CloudFormation template for this.

Option D is invalid because Autoscaling Group will still launch EC2 instances with the older launch configuration

For more information on Autoscaling Launch configuration, please refer to the below document link: from AWS

➤ [http://docs.aws.amazon.com/autoscaling/latest/userguide/l\\_launchConfiguration.html](http://docs.aws.amazon.com/autoscaling/latest/userguide/l_launchConfiguration.html)

#### 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 databases
- 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 SNS
- 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 encryption
- J. Configure the application to force SSL for incoming connections and configure RDS to only grant access if the session is encrypted
- 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 databases
- M. Store RDS credentials in EC2 Parameter Store
- 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 utility
- 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 created
- D. Then build a pipeline in AWS CodePipeline to execute the automation document to build the AMI when triggered
- 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 application
- G. Then start a new EC2 instance from the snapshot and update the running instance using an AWS Lambda function
- H. Take a snapshot of the updated instance, then convert it to an AMI
- I. Store the AMI identification output in an Amazon DynamoDB table.
- J. Launch an Amazon EC2 instance and install Packer
- K. Then configure a Packer build with values defining how the image should be created
- L. Build a Jenkins pipeline to invoke the Packer build when triggered to build an AMI
- M. Store the AMI identification output in an Amazon DynamoDB table.

**Answer: D**

#### NEW QUESTION 10

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 10

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 13

A DevOps Engineer is leading the implementation for automating patching of Windows-based workstations in a hybrid cloud environment by using AWS Systems Manager (SSM).  
What steps should the Engineer follow to set up Systems Manager to automate patching in this environment? (Select TWO.)

- A. Create multiple IAM service roles for Systems Manager so that the ssm.amazonaws.com service can execute the AssumeRole operation on every instance
- B. Register the role on a per-resource level to enable the creation of a service token
- C. Perform managed-instance activation with the newly created service role attached to each managed instance.
- D. Create an IAM service role for Systems Manager so that the ssm.amazonaws.com service can execute the AssumeRole operation
- E. Register the role to enable the creation of a service token



- F. Perform managed-instance activation with the newly created service role.
- G. Using previously obtained activation codes and activation IDs, download and install the SSM Agent on the hybrid servers, and register the servers or virtual machines on the Systems Manager service
- H. Hybrid instances will show with an "mi-" prefix in the SSM console.
- I. Using previously obtained activation codes and activation IDs, download and install the SSM Agent on the hybrid servers, and register the servers or virtual machines on the Systems Manager service
- J. Hybrid instances will show with an "i-" prefix in the SSM console as if they were provisioned as a regular Amazon EC2 instance.
- K. Run AWS Config to create a list of instances that are unpatched and not compliant
- L. Create an instance scheduler job, and through an AWS Lambda function, perform the instance patching to bring them up to compliance.

**Answer:** BC

**Explanation:**

<https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-managed-instance-activation.html>

<https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-install-managed-win.html>

**NEW QUESTION 16**

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

**Answer:** D

**NEW QUESTION 21**

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 cluster.

**Answer:** D

**NEW QUESTION 26**

A company wants to use Amazon DynamoDB for maintaining metadata on its forums. See the sample data set in the image below.

## Thread

| ForumName | Subject | LastPostDateTime      | Thread |
|-----------|---------|-----------------------|--------|
| "S3"      | "aaa"   | "2015-03-15:17:24:31" | 12     |
| "S3"      | "bbb"   | "2015-01-22:23:18:01" | 3      |
| "S3"      | "ccc"   | "2015-02-31:13:14:21" | 4      |
| "S3"      | "ddd"   | "2015-01-03:09:21:11" | 9      |
|           |         |                       |        |
| "EC2"     | "yyy"   | "2015-02-12:11:07:56" | 18     |
| "EC2"     | "zzz"   | "2015-01-18:07:33:42" | 0      |
|           |         |                       |        |
| "RDS"     | "ttt"   | "2015-01-19:01:13:24" | 3      |
| "RDS"     | "sss"   | "2015-03-11:06:53:00" | 11     |
| "RDS"     | "uuu"   | "2015-10-22:12:19:44" | 5      |

A DevOps Engineer is required to define the table schema with the partition key, the sort key, the local secondary index, projected attributes, and fetch operations. The schema should support the following example searches using the least provisioned read capacity units to minimize cost.

- Search within ForumName for items where the subject starts with "a".
- Search forums within the given LastPostDateTime time frame.
- Return the thread value where LastPostDateTime is within the last three months. Which schema meets the requirements?

- A. Use Subject as the primary key and ForumName as the sort key
- B. Have LSI with LastPostDateTime as the sort key and fetch operations for thread.
- C. Use ForumName as the primary key and Subject as the sort key
- D. Have LSI with LastPostDateTime as the sort key and the projected attribute thread.
- E. Use ForumName as the primary key and Subject as the sort key
- F. Have LSI with Thread as the sort key and the projected attribute LastPostDateTime.
- G. Use Subject as the primary key and ForumName as the sort key
- H. Have LSI with Thread as the sort key and fetch operations for LastPostDateTime.

**Answer:** B

### Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/LSI.html>

### NEW QUESTION 30

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 Kinesis
- E. Configure AWS CloudTrail to deliver the API logs to Kinesis
- F. Use Kinesis to load the data into Amazon Redshift
- 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 35

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

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

**Answer:** AC

**NEW QUESTION 38**

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

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

**Answer: B**

**NEW QUESTION 41**

A DevOps engineer notices that all Amazon EC2 instances running behind an Application Load Balancer in an Auto Scaling group are failing to respond to user requests. The EC2 instances are also failing target group HTTP health checks.

Upon inspection, the engineer notices the application process was not running in any EC2 instances. There are a significant number of out of memory messages in the system logs. The engineer needs to improve the resilience of the application to cope with a potential application memory leak. Monitoring and notifications should be enabled to alert when there is an issue.

Which combination of actions will meet these requirements? (Select TWO.)

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

**Answer: BE**

**NEW QUESTION 45**

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

**NEW QUESTION 46**

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

Which of the following should successfully install the application while complying with the new rule?

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

**Answer: C**

**Explanation:**

EC2 instances running in private subnets of a VPC can now have controlled access to S3 buckets, objects, and API functions that are in the same region as the VPC. You can use an S3 bucket policy to indicate which VPCs and which VPC Endpoints have access to your S3 buckets 1-

<https://aws.amazon.com/pt/blogs/aws/new-vpc-endpoint-for-amazon-s3/>



#### NEW QUESTION 48

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 master
- 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 CodeCommit
- 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 group
- G. Create a trigger to notify senior developers with a URL link to approve or deny commit requests delivered through Amazon SNS
- 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 branch
- J. Create separate IAM groups for senior developers and developer
- K. Use an IAM policy to assign each IAM group their corresponding branches
- 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 CodeCommit
- 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 function
- O. Validate the user name with the IAM group to approve or deny the commit.

**Answer: C**

#### NEW QUESTION 53

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 account
- B. Ensure that Access Advisor is configured with a tracking period of 90 days
- C. Invoke the Lambda function using an Amazon CloudWatch Events rule on a weekly schedule
- D. For each record, by user, by service, if the Access Advisor Last Accessed field indicates a day count instead of "Not accessed in the tracking period," this indicates a delta compared to what is in the user's currently attached access policies
- 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 bucket
- G. Create an Amazon Athena table and specify the S3 bucket ARN in the CREATE TABLE query
- 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 eventTime
- 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 week
- 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 account
- M. Ensure that all logs are being sent to a centralized Amazon S3 bucket, so all flow logs can be consolidated and aggregated
- N. Create an AWS Lambda function that is triggered once a week by an Amazon CloudWatch Events schedule
- 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 name
- P. Then have the function determine the deltas on a user-by-user basis
- 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 function
- 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 function
- T. Create the Lambda function to consume the events, parse the input from JSON, and transform it to an Amazon ES document format
- U. POST the documents to the Amazon ES cluster's endpoint by way of the passed-in environment variable
- V. Make sure that the proper indexing exists in Amazon ES and use Apache Lucene queries to parse the permissions on a user-by-user basis
- W. 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 54

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 Firehose stream that has an Amazon S3 bucket destination
- 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 Firehose stream that has an Amazon S3 bucket destination
- 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 Firehose stream that has an Amazon S3 bucket destination
- 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 Firehose stream that has an Amazon S3 bucket destination
- 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 58**

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 activity
- 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 activity
- D. Configure the static website hosting to return the required security headers.
- E. Use an Amazon S3 bucket configured for website hosting
- 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 hosting
- H. Create an Amazon CloudFront distribution that refers to this S3 bucket
- 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 61**

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

How can this be accomplished?

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

**Answer: D**

**NEW QUESTION 63**

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 CloudWatch
- 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 application and database logs into an Amazon S3 bucket
- 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 65**

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

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

**Answer: D**

**Explanation:**

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance

For more information on Autoscaling health checks, please refer to the below document link: from AWS

➤ <http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

**NEW QUESTION 66**

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 traffic configured to route to each stack.
- B. Configure Amazon CloudFront with an AWS Lambda@Edge function to set a cookie when CloudFront receives a request
- 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 request
- 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 request
- 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](https://docs.aws.amazon.com/zh_cn/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html)

#### NEW QUESTION 71

An ecommerce company is looking for ways to deploy an application on AWS that satisfies the following requirements:

- Has a simple and automated application deployment process.
- Has minimal deployment costs while ensuring that at least half of the instances are available to receive end-user requests.
- If the application fails, an automated healing mechanism will replace the affected instances. Which deployment strategy will meet these requirements?

- A. Create an AWS Elastic Beanstalk environment and configure it to use Auto Scaling and an Elastic Load Balance
- B. Use rolling deployments with a batch size of 50%.
- C. Create an AWS OpsWorks stack
- D. Configure the application layer to use rolling deployments as a deployment strategy
- E. Add an Elastic Load Balancing layer
- F. Enable auto healing on the application layer.
- G. Use AWS CodeDeploy with Auto Scaling and an Elastic Load Balancer Use the CodeDeployDefault.HalfAtATime deployment strategy
- H. Enable an Elastic Load Balancing health check to report the status of the application, and set the Auto Scaling health check to ELB.
- I. Use AWS CodeDeploy with Auto Scaling and an Elastic Load Balance
- J. Use a blue/green deployment strategy
- K. Enable an Elastic Load Balancing health check to report the status of the application, and set the Auto Scaling health check to ELB.

**Answer: C**

#### NEW QUESTION 74

A company has developed a Node.js web application which provides REST services to store and retrieve time series data. The web application is built by the Development team on company laptops, tested locally, and manually deployed to a single on-premises server, which accesses a local MySQL database. The company is starting a trial in two weeks, during which the application will undergo frequent updates based on customer feedback. The following requirements must be met:

\*The team must be able to reliably build, test, and deploy new updates on a daily basis, without downtime or degraded performance.

\*The application must be able to scale to meet an unpredictable number of concurrent users during the trial. Which action will allow the team to quickly meet these objectives?

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

**Answer: C**

#### NEW QUESTION 77

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 services
- B. Use AWS CloudFormation 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 services
- 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 78

An e-commerce company is running a web application in an AWS Elastic Beanstalk environment. In recent months, the average load of the Amazon EC2 instances has been increased to handle more traffic.

The company would like to improve the scalability and resilience of the environment. The Development team has been asked to decouple long-running tasks from the environment if the tasks can be executed asynchronously. Examples of these tasks include confirmation emails when users are registered to the platform, and processing images or videos. Also, some of the periodic tasks that are currently running within the web server should be offloaded.

What is the most time-efficient and integrated way to achieve this?

A. Create an Amazon SQS queue and send the tasks that should be decoupled from the Elastic Beanstalk web server environment to the SQS queue

B. Create a fleet of EC2 instances under an Auto Scaling group

C. Use an AMI that contains the application to process the asynchronous tasks, configure the application to listen for messages within the SQS queue, and create periodic tasks by placing those into the cron in the operating system

D. Create an environment variable within the Elastic Beanstalk environment with a value pointing to the SQS queue endpoint.

E. Create a second Elastic Beanstalk worker tier environment and deploy the application to process the asynchronous tasks there

F. Send the tasks that should be decoupled from the original Elastic Beanstalk web server environment to the auto-generated Amazon SQS queue by the Elastic Beanstalk worker environment

G. Place a cron.yaml file within the root of the application source bundle for the worker environment periodic task

H. Use environment links to link the web server environment with the worker environment.

I. Create a second Elastic Beanstalk web server tier environment and deploy the application to process the asynchronous tasks

J. Send the tasks that should be decoupled from the original Elastic Beanstalk web server to the auto-generated Amazon SQS queue by the Elastic Beanstalk web server tier environment

K. Place a cron.yaml file within the root of the application source bundle for the second web server tier environment with the necessary periodic task

L. Use environment links to link both web server environments.

M. Create an Amazon SQS queue and send the tasks that should be decoupled from the Elastic Beanstalk web server environment to the SQS queue

N. Create a fleet of EC2 instances under an Auto Scaling group

O. Install and configure the application to listen for messages within the SQS queue from UserData and create periodic tasks by placing those into the cron in the operating system

P. Create an environment variable within the Elastic Beanstalk web server environment with a value pointing to the SQS queue endpoint.

**Answer: C**

#### NEW QUESTION 81

A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.

Which solution will accomplish this with the LEAST amount of development effort?

A. Create an Amazon CloudWatch Event rule that runs periodically and targets an AWS Lambda function. Within the Lambda function, evaluate the current state of the AWS environment and compare deployed resource values to resource limits on the account

B. Notify the Senior Manager if the account is approaching a service limit.

C. Deploy an AWS Lambda function that refreshes AWS Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically

D. Create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda function

E. In the target Lambda function, notify the Senior Manager.

F. Deploy an AWS Lambda function that refreshes AWS Personal Health Dashboard checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically

G. Create another CloudWatch Events rule with an event pattern matching Personal Health Dashboard events and a target Lambda function

H. In the target Lambda function, notify the Senior Manager.

I. Add an AWS Config custom rule that runs periodically, checks the AWS service limit status, and streams notifications to an Amazon SNS topic

J. Deploy an AWS Lambda function that notifies the Senior Manager, and subscribe the Lambda function to the SNS topic.

**Answer: B**

#### NEW QUESTION 82

A company has thousands of Amazon EC2 instances as well as hundreds of virtual machines on-premises. Developers routinely sign in to the console for on-premises systems to perform troubleshooting. The developers want to sign in to AWS instances to run performance tools, but are unable to due to the lack of a central console logging system. A DevOps engineer wants to ensure that console access is logged on all systems.

Which combination of steps will meet these requirements? (Select TWO.)

A. Attach a role to all AWS instances that contains the appropriate permission

B. Create an AWS Systems Manager managed-instance activation

C. Install and configure Systems Manager Agent on on-premises machines.

D. Enable AWS Systems Manager Session Manager logging to an Amazon S3 bucket

E. Direct developers to connect to the systems with Session Manager only.

F. Enable AWS Systems Manager Session Manager logging to AWS CloudTrail

G. Direct developers to continue normal sign-in procedures for on-premise

H. Use Session Manager for AWS instances.

I. Install and configure an Amazon CloudWatch Logs agent on all systems

J. Create an AWS Systems Manager managed-instance activation.

K. Set up a Site-to-Site VPN connection between the on-premises and AWS network

L. Set up a bastion instance to allow developers to sign in to the AWS instances.



**Answer:** AB

#### NEW QUESTION 83

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 cron
- 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 server
- 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 88

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

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

**Answer:** D

#### NEW QUESTION 91

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 94

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

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

**Answer:** D

#### Explanation:

If you suspend AddToLoadBalancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddToLoadBalancer process, Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does not add the instances that were launched while this process was suspended. You must register those instances manually.

Option A is invalid because this just balances the number of EC2 instances in the group across the Availability Zones in the region

Option B is invalid because this just checks the health of the instances. Auto Scaling marks an instance as unhealthy if Amazon EC2 or Elastic Load Balancing tells

Auto Scaling that the instance is unhealthy.

Option C is invalid because this process just terminates instances that are marked as unhealthy and later creates new instances to replace them.

For more information on process suspension, please refer to the below document link: from AWS

➤ <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

#### NEW QUESTION 95

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 with weighted 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 code
- E. Swap the Elastic Beanstalk environment CNAME
- 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 stack
- 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 validated
- L. Keep the old production environment in place until the new staging environment handles all traffic.

**Answer:** D

#### NEW QUESTION 97

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 activities
- B. Run backup scripts on Amazon EC2 in an Auto Scaling group
- 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 activities
- E. Run backup scripts on Amazon EC2 in an Auto Scaling group
- F. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading backup logs to Amazon Redshift
- 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 activities
- I. Run backup scripts on Amazon EC2 in a single Availability Zone
- J. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading the backup logs to Amazon RDS
- 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 activities
- M. Run backup scripts on Amazon EC2 in a single Availability Zone
- 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 administrators with backup activity metadata.

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 99

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

Which deployment sequence will meet these requirements?

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

**Answer:** A

#### NEW QUESTION 101

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

Which is the MOST cost-effective solution?

- A. Create an EC2 Instance-terminate Lifecycle Action on the group, write a termination script for pushing logs into Amazon S3, and trigger an AWS Lambda function based on S3 PUT to create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance

Termination Date.

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

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 103

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 107

A company is deploying a container-based application using AWS CodeBuild. The security team mandates that all containers are scanned for vulnerabilities prior to deployment using a password-protected endpoint. All sensitive information must be stored securely.

Which solution should be used to meet these requirements?

- A. Encrypt the password using AWS KM
- B. Store the encrypted password in the buildspec.yml file as an environment variable under the variables mappin
- C. Reference the environment variable to initiate scanning.
- D. Import the password into an AWS CloudHSM ke
- E. Reference the CloudHSM key in the buildpec.yml file as an environment variable under the variables mappin
- F. Reference the environment variable to initiate scanning.
- G. Store the password in the AWS Systems Manager Parameter Store as a secure strin
- H. Add the Parameter Store key to the buildspec.yml file as an environment variable under the parameter-store mappin
- I. Reference the environment variable to initiate scanning.
- J. Use the AWS Encryption SDK to encrypt the password and embed in the buildspec.yml file as a variable under the secrets mappin
- K. Attach a policy to CodeBuild to enable access to the required decryption key.

**Answer:** C

#### NEW QUESTION 108

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 running during the update? Which of the options below listed can be used for this?

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

**Answer:** A

**Explanation:**

The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scaling group 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 110

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 and
- 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 113

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 116

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

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

**Answer:** A

#### NEW QUESTION 118

A company is running an application on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones.

After a recent application update, users are getting HTTP 502 Bad Gateway errors from the application URL. The DevOps Engineer cannot analyze the problem because Auto Scaling is terminating all EC2 instances shortly after launch for being unhealthy.

What steps will allow the DevOps Engineer access to one of the unhealthy instances to troubleshoot the deployed application?

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 123

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 repository
- 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 repository
- 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 125

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

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

**Answer:** A

#### NEW QUESTION 128

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

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

**Answer:** C

#### NEW QUESTION 129

A company runs an application with an Amazon EC2 and on-premises configuration. A DevOps engineer needs to standardize patching across both environments. Company policy dictates that patching only happens during non-business hours. Which combination of actions will meet these requirements? (Select THREE.)

- A. Add the physical machines into AWS Systems Manager using Systems Manager Hybrid Activations.
- B. Attach an IAM role to the EC2 instances, allowing them to be managed by AWS Systems Manager.
- C. Create IAM access keys for the on-premises machines to interact with AWS Systems Manager.
- D. Execute an AWS Systems Manager Automation document to patch the systems every hour.
- E. Use Amazon CloudWatch Events scheduled events to schedule a patch window.
- F. Use AWS Systems Manager Maintenance Windows to schedule a patch window.

**Answer:** ABF

#### NEW QUESTION 132

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

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

**Answer:** A

**NEW QUESTION 137**

After presenting a working proof of concept for a new application that uses AWS API Gateway, a Developer must set up a team development environment for the project. Due to a tight timeline, the Developer wants to minimize time spent on infrastructure setup, and would like to reuse the code repository created for the proof of concept. Currently, all source code is stored in AWS CodeCommit.

Company policy mandates having alpha, beta, and production stages with separate Jenkins servers to build code and run tests for every stage. The Development Manager must have the ability to block code propagation between admins at any time. The Security team wants to make sure that users will not be able to modify the environment without permission.

How can this be accomplished?

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

**Answer: D**

**NEW QUESTION 142**

A company is implementing a well-architected design for its globally accessible API stack. The design needs to ensure both high reliability and fast response times for users located in North America and Europe.

The API stack contains the following three tiers:

- Amazon API Gateway
- AWS Lambda
- Amazon DynamoDB

Which solution will meet the requirements?

- A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using health check
- B. Configure the APIs to forward requests to a Lambda function in that Regio
- C. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.
- D. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using latency-based routing and health check
- E. Configure the APIs to forward requests to a Lambda function in that Regio
- F. Configure the Lambda functions to retrieve and update the data in a DynamoDB global table.
- G. Configure Amazon Route 53 to point to API Gateway in North America, create a disaster recovery API in Europe, and configure both APIs to forward requests to the Lambda functions in that Regio
- H. Retrieve the data from a DynamoDB global tabl
- I. Deploy a Lambda function to check the North America API health every 5 minute
- J. In the event of a failure, update Route 53 to point to the disaster recovery API.
- K. Configure Amazon Route 53 to point to API Gateway API in North America using latency-based routin
- L. Configure the API to forward requests to the Lambda function in the Region nearest to the use
- M. Configure the Lambda function to retrieve and update the data in a DynamoDB table.

**Answer: B**

**NEW QUESTION 144**

An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and maintained with AWS Cloud Formation. A DevOps engineer wants the instances to have the latest configuration file when launched, and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files in source control.

Which solution will accomplish this?

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

**Answer: B**

**NEW QUESTION 146**

A company is creating a software solution that executes a specific parallel-processing mechanism. The software can scale to tens of servers in some special scenarios. This solution uses a proprietary library that is license-based, requiring that each individual server have a single, dedicated license installed. The company has 200 licenses and is planning to run 200 server nodes concurrently at most.

The company has requested the following features:

"ç A mechanism to automate the use of the licenses at scale. "ç Creation of a dashboard to use in the future to verify which licenses are available at any moment.

What is the MOST effective way to accomplish these requirements?

- A. Upload the licenses to a private Amazon S3 bucke
- B. Create an AWS CloudFormation template with a Mappings section for the license
- C. In the template, create an Auto Scaling group to launch the server



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

**Answer:** D

#### NEW QUESTION 149

An Amazon EC2 instance with no internet access is running in a Virtual Private Cloud (VPC) and needs to download an object from a restricted Amazon S3 bucket. When the DevOps Engineer tries to gain access to the object, an Access Denied error is received.

What are the possible causes for this error? (Select THREE.)

- A. The S3 bucket default encryption is enabled.
- B. There is an error in the S3 bucket policy.
- C. There is an error in the VPC endpoint policy.
- D. The object has been moved to Amazon Glacier.
- E. There is an error in the IAM role configuration.
- F. S3 versioning is enabled.

**Answer:** BCE

#### NEW QUESTION 154

You currently have the following setup in AWS

- 1) An Elastic Load Balancer
- 2) Auto Scaling Group which launches EC2 Instances
- 3) AMIs with your code pre-installed

You want to deploy the updates of your app to only a certain number of users. You want to have a cost-effective solution. You should also be able to revert back quickly. Which of the below solutions is the most feasible one?

- A. Create a second ELB, and a new Auto Scaling Group assigned a new Launch Configuration
- B. Create a new AMI with the updated app
- C. Use Route53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs.
- D. Create new AMIs with the new app
- E. Then use the new EC2 instances in half proportion to the older instances.
- F. Redeploy with AWS Elastic Beanstalk and Elastic Beanstalk version
- G. Use Route 53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs
- H. Create a full second stack of instances, cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.

**Answer:** A

#### Explanation:

The Weighted Routing policy of Route53 can be used to direct a proportion of traffic to your application. The best option is to create a second CLB, attach the new Auto Scaling Group and then use Route53 to divert the traffic.

Option B is wrong because just having EC2 instances running with the new code will not help.

Option C is wrong because Elastic Beanstalk is good for development environments, and also there is no mention of having 2 environments where environment URLs can be swapped.

Option D is wrong because you still need Route53 to split the traffic.

For more information on Route53 routing policies, please refer to the below link:

➤ <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

#### NEW QUESTION 156

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

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

**Answer:** D

#### Explanation:

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues.

Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.

For more information on CLB access logs, please refer to the below document link: from AWS

➤ <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html>

#### NEW QUESTION 161

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

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

**Answer:** AC

#### NEW QUESTION 162

Management has reported an increase in the monthly bill from Amazon Web Services, and they are extremely concerned with this increased cost. Management has asked you to determine the exact cause of this increase. After reviewing the billing report, you notice an increase in the data transfer cost. How can you provide management with a better insight into data transfer use?

- A. Update your Amazon CloudWatch metrics to use five-second granularity, which will give better detailed metrics that can be combined with your billing data to pinpoint anomalies.
- B. Use Amazon CloudWatch Logs to run a map-reduce on your logs to determine high usage and data transfer.
- C. Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points.
- D- Using Amazon CloudWatch metrics, pull your Elastic Load Balancing outbound data transfer metrics monthly, and include them with your billing report to show which application is causing higher bandwidth usage

**Answer:** C

#### Explanation:

You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console.

CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a statistic set.

If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.

Option A won't be sufficient to provide better insights.

Option B is an overhead when you can make the application publish custom metrics Option D is invalid because just the ELB metrics will not give the entire picture

For more information on custom metrics, please refer to the below document link: from AWS

➤ <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

#### NEW QUESTION 164

An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages while it is running.

All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.

How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

- A. Add Deletion Policy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.
- B. Add a custom resource when an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM role
- C. Write the Lambda function to delete all objects from the bucket when the RequestType is Delete.
- D. Identify the resource that was not delete
- E. From the S3 console, empty the S3 bucket and then delete it.
- F. Replace the EC2 and S3 bucket resources with a single AWS OpsWorks Stacks resource
- G. Define a custom recipe for the stack to create and delete the EC2 instance and the S3 bucket.

**Answer:** C

#### NEW QUESTION 167

A company maintains a stateless web application that is experiencing inconsistent traffic. The company uses AWS CloudFormation to deploy the application. The application runs on Amazon EC2 On-Demand Instances behind an Application Load Balancer (ALB). The instances run across multiple Availability Zones.

The company wants to include the use of Spot Instances while continuing to use a small number of On-Demand Instances to ensure that the application remains highly available.

What is the MOST cost-effective solution that meets these requirements?

- A. Add a Spot block resource to the AWS CloudFormation template
- B. Use the diversified allocation strategy with step scaling behind the ALB.
- C. Add a Spot block resource to the AWS CloudFormation template
- D. Use the lowest-price allocation strategy with target tracking scaling behind the ALB.
- E. Add a Spot Fleet resource to the AWS CloudFormation template
- F. Use the capacity-optimized allocation strategy with step scaling behind the ALB.
- G. Add a Spot Fleet resource to the AWS CloudFormation template
- H. Use the diversified allocation strategy with scheduled scaling behind the ALB

Answer: C

#### NEW QUESTION 168

An Application team is refactoring one of its internal tools to run in AWS instead of on-premises hardware. All of the code is currently written in Python and is standalone. There is also no external state store or relational database to be queried.

Which deployment pipeline incurs the LEAST amount of changes between development and production?

- A. Developers should use Docker for local developmen
- B. Use AWS SMS to import these containers as AMIs for Amazon EC2 whenever dependencies are update
- C. Use AWS CodePipeline to test new code changes against the Auto Scaling group.
- D. Developers should use their native Python environmen
- E. When Dependencies are changed and a new container is ready, use AWS CodePipeline and AWS CodeBuild to perform functional tests and then upload the new container to the Amazon EC
- F. Use AWS CloudFormation with the custom container to deploy the new Amazon ECS.
- G. Developers should use their native Python environmen
- H. When Dependencies are changed and a new code is ready, use AWS CodePipeline and AWS CodeBuild to perform functional tests and then upload the new container to the Amazon EC
- I. Use CodePipeline and CodeBuild with the custom container to test new code changes inside AWS Elastic Beanstalk

Answer: A

#### NEW QUESTION 169

A company has microservices running in AWS Lambda that read data from Amazon DynamoDB. The Lambda code is manually deployed by Developers after successful testing. The company now needs the tests and deployments be automated and run in the cloud. Additionally, traffic to the new versions of each microservice should be incrementally shifted over time after deployment.

What solution meets all the requirements, ensuring the MOST developer velocity?

- A. Create an AWS CodePipeline configuration and set up a post-commit hook to trigger the pipeline after tests have passe
- B. Use AWS CodeDeploy and create a Canary deployment configuration that specifies the percentage of traffic and interval.
- C. Create an AWS CodeBuild configuration that triggers when the test code is pushe
- D. Use AWS CloudFormation to trigger an AWS CodePipeline configuration that deploys the new Lambda versions and specifies the traffic shift percentage and interval.
- E. Create an AWS CodePipeline configuration and set up the source code step to trigger when code ispushe
- F. Set up the build step to use AWS CodeBuild to run the test
- G. Set up an AWS CodeDeploy configuration to deploy, then select the CodeDeployDefault.LambdaLinear10PercentEvery3Minutes option.
- H. Use the AWS CLI to set up a post-commit hook that uploads the code to an Amazon S3 bucket after tests have passe
- I. Set up an S3 event trigger that runs a Lambda function that deploys the new versio
- J. Use an interval in the Lambda function to deploy the code over time at the required percentage.

Answer: C

#### Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html>

#### NEW QUESTION 173

A DevOps Engineer has been asked by the Security team to ensure that AWS CloudTrail files are not tampered with after being created. Currently, there is a process with multiple trails, using AWS IAM to restrict access to specific trails. The Security team wants to ensure they can trace the integrity of each file and make sure there has been no tampering.

Which option will require the LEAST effort to implement and ensure the legitimacy of the file while allowing the Security team to prove the authenticity of the logs?

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

Answer: C

#### Explanation:

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

#### NEW QUESTION 174

.....



## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-DevOps-Engineer-Professional Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-DevOps-Engineer-Professional Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-DevOps-Engineer-Professional/>

## Money Back Guarantee

### **AWS-Certified-DevOps-Engineer-Professional Practice Exam Features:**

- \* AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently
- \* AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff
- \* AWS-Certified-DevOps-Engineer-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year