Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1
A company has migrated its container-based applications to Amazon EKS and want to establish automated email notifications. The notifications sent to each email address are for specific activities related to EXS components. The solution will include Amazon SNS topics and an AWS Lambda function to evaluate incoming log events and publish messages to the correct SNS topic.
Which logging solution will support these requirements?

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

Answer: A

NEW QUESTION 2
An application runs on Amazon EC2 instances behind an Application Load Balancer. Amazon RDS MySOL is used on the backend. The instances run in an Auto Scaling group across multiple Availability Zones. The Application Load Balancer health check ensures the web servers are operating and able to make read/write SQL connections. Amazon Route 53 provides DNS functionality with a record pointing to the Application Load Balancer. A new policy requires a geographically isolated disaster recovery site with an RTO of 4 hours and an RPO of 15 minutes.
Which disaster recovery strategy will require the LEAST amount of changes to the application stack?

A. Launch a replica stack of everything except RDS in a different Availability Zon
B. Create an RDS read-only replica in a new Availability Zone and configure the new stack to point to the local RDS instanc
C. Add the new stack to the Route 53 record set with a failover routing policy.
D. Launch a replica stack of everything except RDS in a different regio
E. Create an RDS read-only replica in a new region and configure the new stack to point to the local RDS instanc
F. Add the new stack to the Route 53 record set with a latency routing policy.
G. Launch a replica stack of everything except RDS in a different regio
H. Upon failure, copy the snapshot over from the primary region to the disaster recovery regio
I. Adjust the Amazon Route 53 record set to point to the disaster recovery region's Application Load Balancer.
J. Launch a replica stack of everything except RDS in a different regio
K. Create an RDS read-only replica in a new region and configure the new stack to point to the local RDS instanc
L. Add the new stack to the Amazon Route 53 record set with a failover routing policy

Answer: D

NEW QUESTION 3
A company wants 10 use AWS development tools to replace Its current bash deployment scripts. The company currently deploys a LAMP application to a group of Amazon EC2 instances behind an Application Load Balancer (ALB). During the deployments, the company unit tests the committed application, stops and starts services, unregisters and re-registers instances with the load balancer, and updates Me permissions. The company wants to maintain the same deployment functionality through the shift to using AWS services.
Which solution will meet these requirements?

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

Answer: D

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

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

Answer: D

NEW QUESTION 5
A DevOps Engineer has a single Amazon DynamoDB table that received shipping orders and tracks inventory. The Engineer has three AWS Lambda functions reading from a DymamoDB stream on that table. The Lambda functions perform various functions such as doing an item count, moving items to Amazon Kinesis Data Firehose, monitoring inventory levels, and creating vendor orders when parts are low.
While reviewing logs, the Engineer notices the Lambda functions occasionally fail under increased load, receiving a stream throttling error.
Which is the MOST cost-effective solution that requires the LEAST amount of operational management?

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

Answer: C

NEW QUESTION 6
A company that uses electronic health records is running a fleet of Amazon EC2 instances with an Amazon Linux operating system. As part of patient privacy requirements, the company must ensure continuous compliance for patches for operating system and applications running on the EC2 instances.
How can the deployments of the operating system and application patches be automated using a default and custom repository?

A. Use AWS Systems Manager to create a new patch baseline including the custom repositor
B. Execute the AWS-RunPatchBaseline document using the run command to verify and install patches.
C. Use AWS Direct Connect to integrate the corporate repository and deploy the patches using Amazon CloudWatch scheduled events, then use the CloudWatch dashboard to create reports.
D. Use yum-config-manager to add the custom repository under /etc/yum.repos.d and run yum-config-manager-enable to activate the repository.
E. Use AWS Systems Manager to create a new patch baseline including the corporate repositor
F. Execute the AWS-AmazonLinuxDefaultPatchBaseline document using the run command to verify and install patches.

Answer: A

Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-about-aws-runpatchbaseline.htm

NEW QUESTION 7
A mobile application running on eight Amazon EC2 instances is relying on a third-party API endpoint. The thirdparty service has a high failure rate because of limited capacity, which is expected to be resolved in a few weeks. In the meantime, the mobile application developers have added a retry mechanism and are logging failed API requests. A DevOps Engineer must automate the monitoring of application logs and count the specific error messages; if there are more than 10 errors within a 1-minute window, the system must issue an alert. How can the requirements be met with MINIMAL management overhead?

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

Answer: C

NEW QUESTION 8
A company has deployed several applications globally. Recently, Security Auditors found that few Amazon EC2 instances were launched without Amazon EBS disk encryption. The Auditors have requested a report detailing all EBS volumes that were not encrypted in multiple AWS accounts and regions. They also want to be notified whenever this occurs in future.
How can this be automated with the LEAST amount of operational overhead?

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

Answer: C

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

NEW QUESTION 9
Your company has multiple applications running on AWS. Your company wants to develop a tool that notifies on-call teams immediately via email when an alarm is triggered in your environment. You have multiple on-call teams that work different shifts, and the tool should handle notifying the correct teams at the correct times.
How should you implement this solution?

A. Create an Amazon SNS topic and an Amazon SQS queu
B. Configure the Amazon SQS queue as a subscriber to the Amazon SNS topic.Configure CloudWatch alarms to notify this topic when an alarm is triggere
C. Create an Amazon EC2 Auto Scaling group with both minimum and desired Instances configured to 0. Worker nodes in thisgroup spawn when messages are added to the queu
D. Workers then use Amazon Simple Email Service to send messages to your on call teams.
E. Create an Amazon SNS topic and configure your on-call team email addresses as subscriber
F. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to this new topi
G. Notifications will be sent to on-call users when a CloudWatch alarm is triggered.
H. Create an Amazon SNS topic and configure your on-call team email addresses as subscriber
I. Create a secondary Amazon SNS topic for alarms and configure your CloudWatch alarms to notify this topic when triggere
J. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggere
K. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the first topic so that on-call engineers receive alerts.
L. Create an Amazon SNS topic for each on-call group, and configure each of these with the team member emails as subscriber
M. Create another Amazon SNS topic and configure your CloudWatch alarms to notify this topic when triggere
N. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggere
O. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the correct team topic when on shift.

Answer: D

Explanation:
Option D fulfils all the requirements
1) First is to create a SNS topic for each group so that the required members get the email addresses.
2) Ensure the application uses the HTTPS endpoint and the SDK to publish messages Option A is invalid because the SQS service is not required.
Option B and C are incorrect. As per the requirement we need to provide notification to only those on-call teams who are working in that particular shift when an alarm is triggered. It need not have to be send to all the on-call teams of the company. With Option B & C, since we are not configuring the SNS topic for each on call team the notifications will be send to all the on-call teams. Hence these 2 options are invalid. For more information on setting up notifications, please refer to the below document link: from AWS
http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

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 10
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 instanc
B. Register the role on a per-resource level to enable the creation of a service toke
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 operatio
E. Register the role to enable the creation of a service toke
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 servic
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 servic
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 complian
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 13
A company has an application that has predictable peak traffic times. The company wants the application instances to scale up only during the peak times. The application stores state in Amazon DynamoDB. The application environment uses a standard Node.js application stack and custom Chef recipes stored in a private Git repository.
Which solution is MOST cost-effective and requires the LEAST amount of management overhead when performing rolling updates of the application environment?

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

Answer: D

NEW QUESTION 16
A Development team creates a build project in AWS CodeBuild. The build project invokes automated tests of modules that access AWS services.
Which of the following will enable the tests to run the MOST securely?

A. Generate credentials for an IAM user with a policy attached to allow the actions on AWS service
B. Store credentials as encrypted environment variables for the build projec
C. As part of the build script, obtain the credentials to run the integration tests.
D. Have CodeBuild run only the integration tests as a build job on a Jenkins serve
E. Create a role that has a policy attached to allow the actions on AWS service
F. Generate credentials for an IAM user that is allowed to assume the rol
G. Configure the credentials as secrets in Jenkins, and allow the build job to use them to run the integration tests.
H. Create a service role in IAM to be assumed by CodeBuild with a policy attached to allow the actions on AWS service
I. Configure the build project to use the role created.
J. Use AWS managed credential
K. Encrypt the credentials with AWS KM
L. As part of the build script, decrypt with AWS KMS and use these credentials to run the integration tests.

Answer: B

NEW QUESTION 21
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"	"rrr"	"2015-01-19:01:13:24"	3
"RDS"	"sss"	"2015-03-11:06:53:00"	11
"RDS"	"ttt"	"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 ke
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 ke
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 ke
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 ke
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 22
A company is deploying a new application that uses Amazon EC2 instances. The company needs a solution to query application logs and AWS account API activity. Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 27
An application has microservices spread across different AWS accounts and is integrated with an on-premises legacy system for some of its functionality. Because of the segmented architecture and missing logs, every time the application experiences issues, it is taking too long to gather the logs to identify the issues. A DevOps Engineer must fix the log aggregation process and provide a way to centrally analyze the logs.
Which is the MOST efficient and cost-effective solution?

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

Answer: D

NEW QUESTION 32
A company is running an application on Amazon EC2 instances in an Auto Scaling group. Recently, an issue occurred that prevented EC2 instances from launching successfully, and it took several hours for the support team to discover the issue. The support team wants to be notified by email whenever an EC2 instance does not start successfully.
Which action will accomplish this?

A. Add a health check to the Auto Scaling group to invoke an AWS Lambda function whenever an instance status is impaired.
B. Configure the Auto Scaling group to send a notification to an Amazon SNS topic whenever a failed instance launch occurs.
C. Create an Amazon CloudWatch alarm that invokes an AWS Lambda function when a failed AttachInstances Auto Scaling API call is made.
D. Create a status check alarm on Amazon EC2 to send a notification to an Amazon SNS topic whenever a status check fail occurs.

Answer: B

NEW QUESTION 37
A DevOps engineer is architecting a continuous development strategy for a company's software as a service (SaaS) web application running on AWS. For application and security reasons, users subscribing to this application are distributed across multiple Application Load Balancers (ALBs), each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances. The application does not require a build stage, and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs. Auto Scaling groups, and EC2 fleets.
Which architecture will meet these requirements with the LEAST amount of configuration?

A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair.
B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.
C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.
D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and

deployment group created for the same ALB-Auto Scaling group pair.

Answer: C

NEW QUESTION 39
A DevOps Engineer encountered the following error when attempting to use an AWS CloudFormation template to create an Amazon ECS cluster:
An error occurred (InsufficientCapabilitiesException) when calling the CreateStack operation.
What caused this error and what steps need to be taken to allow the Engineer to successfully execute the AWS CloudFormation template?

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

Answer: C

NEW QUESTION 42
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 47
Your application stores sensitive information on an EBS volume attached to your EC2 instance. How can you protect your information? Choose two answers from the options given below

A. Unmount the EBS volume, take a snapshot and encrypt the snapsho
B. Re-mount the Amazon EBS volume
C. It is not possible to encrypt an EBS volume, you must use a lifecycle policy to transfer data to S3 for encryption.
D. Copy the unencrypted snapshot and check the box to encrypt the new snapsho
E. Volumes restored from this encrypted snapshot will also be encrypted.
F. Create and mount a new, encrypted Amazon EBS volum
G. Move the data to the new volum
H. Delete the old Amazon EBS volume *t

Answer: CD

Explanation:
These steps are given in the AWS documentation
To migrate data between encrypted and unencrypted volumes
1) Create your destination volume (encrypted or unencrypted, depending on your need).
2) Attach the destination volume to the instance that hosts the data to migrate.
3) Make the destination volume available by following the procedures in Making an Amazon EBS Volume Available for Use. For Linux instances, you can create a mount point at /mnt/destination and mount the destination volume there.
4) Copy the data from your source directory to the destination volume. It may be most convenient to use a bulk-copy utility for this.
To encrypt a volume's data by means of snapshot copying
1) Create a snapshot of your unencrypted CBS volume. This snapshot is also unencrypted.
2) Copy the snapshot while applying encryption parameters. The resulting target snapshot is encrypted.
3) Restore the encrypted snapshot to a new volume, which is also encrypted.
For more information on EBS Encryption, please refer to the below document link: from AWS
> http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ EBSEncryption.html

NEW QUESTION 51
A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency.
Which actions should be taken to accomplish this? {Select TWO.)

A. Install the CloudWatch agent server side and configure the agent to upload relevant logs to CloudWatch.
B. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and upload those segments to X-Ray during each request.
C. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and use the X-Ray daemon to upload segments to X-Ray.

D. Modify the on-premises application to send log information back to API Gateway with each request.
E. Modify the on-premises application to calculate and upload statistical data relevant to the API service requests to CloudWatch metrics.

Answer: AC

NEW QUESTION 52
A company is using AWS CodeDeploy to manage its application deployments. Recently, the Development team decided to use GitHub for version control, and the team is looking for ways to integrate the GitHub repository with CodeDeploy. The team also needs to develop a way to automate deployment whenever there is a new commit on that repository. The team is currently deploying new application revisions by manually indicating the Amazon S3 location.
How can the integration be achieved in the MOST efficient way?

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

Answer: B

NEW QUESTION 56
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 60
A company is migrating an application to AWS that runs on a single Amazon EC2 instance. Because of licensing limitations, the application does not support horizontal scaling. The application will be using Amazon Aurora for its database.
How can the DevOps Engineer architect automated healing to automatically recover from EC2 and Aurora failures, in addition to recovering across Availability Zones (AZs), in the MOST cost-effective manner?

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

Answer: C

NEW QUESTION 62
A government agency has multiple AWS accounts, many of which store sensitive citizen information. A Security team wants to detect anomalous account and network activities (such as SSH brute force attacks) in any account and centralize that information in a dedicated security account. Event information should be stored in an Amazon S3 bucket in the security account, which is monitored by the department's Security Information and Even Manager (SIEM) system.
How can this be accomplished?

A. Enable Amazon Macie in every accoun
B. Configure the security account as the Macie Administrator for every member account using invitation/acceptanc
C. Create an Amazon CloudWatch Events rule in the security account to send all findings to Amazon Kinesis Data Firehouse, which should push the findings to the

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

Answer: C

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

NEW QUESTION 67
A company uses federated access for its AWS environment The available roles are created and managed using AWS CloudFormation from a CI/CD pipeline. All changes should be made to the IAM roles through the pipeline. The security team found that changes are being made to the roles out-of-band and would like to detect when this occurs.
Which action will accomplish this?

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

Answer: D

NEW QUESTION 68
A healthcare provider has a hybrid architecture that includes 120 on-premises VMware servers running RedHat and 50 Amazon EC2 instances running Amazon Linux. The company is in the middle of an all-in migration to AWS and wants to implement a solution for collecting information from the on-premises virtual machines and the EC2 instances for data analysis. The information includes:
- Operating system type and version
- Data for installed applications
- Network configuration information, such as MAC and IP addresses
- Amazon EC2 instance AMI ID and IAM profile
How can these requirements be met with the LEAST amount of administration?

A. Write a shell script to run as a cron job on EC2 instances to collect and push the data to Amazon S3. For on-premises resources, use VMware vSphere to collect the data and write it into a file gateway forstoring the data in S3. Finally, use Amazon Athena on the S3 bucket for analytics.
B. Use a script on the on-premises virtual machines as well as the EC2 instances to gather and push the data into Amazon S3, and then use Amazon Athena for analytics.
C. Install AWS Systems Manager agents on both the on-premises virtual machines and the EC2 instances.Enable inventory collection and configure resource data sync to an Amazon S3 bucket to analyze the data with Amazon Athena.
D. Use AWS Application Discovery Service for deploying Agentless Discovery Connector in the VMware environment and Discovery Agents on the EC2 instances for collecting the dat
E. Then use the AWS Migration Hub Dashboard for analytics.

Answer: C

NEW QUESTION 71
A company is using AWS CodeCommit as its source code repository. After an internal audit, the compliance team mandates that any code change that go into the master branch must be committed by senior developers.
Which solution will meet these requirements?

A. Create two repositories in CodeCommit: one for working and another for the maste
B. Create separate IAM groups for senior developers and developer
C. Assign the resource-level permissions on the repositories tied to the IAM group
D. After the code changes are reviewed, sync the approved files to the master code commit repository.
E. Create a repository in CodeCommi
F. Create separate IAM groups for senior developers and developers.Assign code commit permissions for both groups, with code merge permissions for the senior developers grou
G. Create a trigger to notify senior developers with a URL link to approve or deny commit requests delivered through Amazon SN
H. Once a senior developer approves the code, the code gets merged to the master branch.
I. Create a repository in CodeCommit with a working and master branc
J. Create separate IAM groups for senior developers and developer
K. Use an IAM policy to assign each IAM group their corresponding branche
L. Once the code is merged to the working branch, senior developers can pull the changes from the working branch to the master branch.
M. Create a repository in CodeCommi
N. Create separate IAM groups for senior developers and developers.Use AWS Lambda triggers on the master branch and get the user name of the developer at the event object of the Lambda functio
O. Validate the user name with the IAM group to approve or deny the commit.

Answer: C

NEW QUESTION 76

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

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

Answer: C

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

NEW QUESTION 80
A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project.
How can this issue be corrected in the MOST secure manner?

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

Answer: C

NEW QUESTION 85
After a recent audit, a company decided to implement a new disaster recovery strategy for Its Amazon S3 data and its MySQL database running on Amazon EC2. Management wants the ability to recover to a secondary AWS Region with an RPO under 5 seconds and a RTO under 1 minute.
Which actions will meet the requirements while MINIMIZING operational overhead? (Select TWO.)

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

Answer: AE

NEW QUESTION 86
A DevOps Engineer must implement monitoring for a workload running on Amazon EC2 and Amazon RDS MySQL. The monitoring must include:
Application logs and operating system metrics for the Amazon EC2 instances Database logs and operating system metrics for the Amazon RDS database Which steps should the Engineer take?

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

Answer: B

NEW QUESTION 91
A company's application is running on Amazon EC2 instances in an Auto Scaling group. A DevOps engineer needs to ensure there are at least four application servers running at all times. Whenever an update has to be made to the application, the engineer creates a new AMI with the updated configuration and updates

the AWS CloudFormation template with the new AMI ID. After the stack update finishes, the engineer manually terminates the old instances one by one. verifying that the new instance is operational before proceeding. The engineer needs to automate this process.
Which action will allow for the LEAST number of manual steps moving forward?

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

Answer: A

NEW QUESTION 94
An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.
When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.
How should the company meet these requirements with the LEAST amount of application changes?

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

Answer: C

NEW QUESTION 95
A DevOps engineer has automated a web service deployment using AWS CodePipelme with the following steps:
• An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
• An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
• A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment The quality assurance (QA) team has asked for permission to inspect the build artifact before the deployment to the production environment occurs. The OA team wants to run an internal automated penetration testing tool (invoked using a REST API call) to run some manual tests.
Which combination of actions will fulfill this request? (Select TWO.)

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

Answer: BD

NEW QUESTION 98
A company wants to use AWS Systems Manager documents to bootstrap physical laptops for developers. The bootstrap code is stored in GitHub. A DevOps engineer has already created a Systems Manager activation, installed the Systems Manager agent with the registration code, and installed an activation ID on all the laptops.
Which set of steps should be taken next?

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

Answer: D

NEW QUESTION 99
A media customer has several thousand amazon EC2 instances in an AWS account. The customer is using a Slack channel for team communications and important updates. A DevOps Engineer was told to send all AWS-scheduled EC2 maintenance notifications to the company Slack channel.
Which method should the Engineer use to implement this process in the LEAST amount of steps?

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

Answer: B

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

NEW QUESTION 101
An education company has a Docker-based application running on multiple Amazon EC2 instances in an Amazon ECS cluster. When deploying a new version of

the application, the Developer, pushes a new image to a private Docker container registry, and then stops and starts all tasks to ensure that they all have the latest version of the application. The Developer discovers that the new tasks are occasionally running with an old image.
How can this issue be prevented?

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

Answer: C

NEW QUESTION 105
The management team at a company with a large on-premises OpenStack environment wants to move non-production workloads to AWS. An AWS Direct Connect connection has been provisioned and configured to connect the environments. Due to contractual obligations, the production workloads must remain on-premises, and will be moved to AWS after the next contract negotiation. The company follows Center for Internet Security (CIS) standards for hardening images; this configuration was developed using the company's configuration management system.
Which solution will automatically create an identical image in the AWS environment without significant overhead?

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

Answer: D

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

NEW QUESTION 110
A company wants to implement a CI/CD pipeline for an application that is deployed on AWS. The company also has a source-code analysis tool hosted on premises that checks for security flaws. The tool has not yet been migrated to AWS and can be accessed only on premises. The company wants to run checks against the source code as part of the pipeline before the code is compiled. The checks take anywhere from minutes to an hour to complete.
How can a DevOps Engineer meet these requirements?

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

Answer: B

NEW QUESTION 113
A company wants to use AWS CloudFormation for infrastructure deployment. The company has strict tagging and resource requirements and wants to limit the deployment to two Regions. Developers will need to deploy multiple versions of the same application.
Which solution ensures resources are deployed in accordance with company policy?

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

Answer: C

NEW QUESTION 114
A company uses AWS CodePipeline to manage and deploy infrastructure as code. The infrastructure is defined in AWS CloudFormation templates and is primarily comprised of multiple Amazon EC2 instances and Amazon RDS databases. The Security team has observed many operators creating inbound security group rules with a source CIDR of 0 0 0 0/0 and would like to proactively stop the deployment of rules with open CIDRs
The DevOps Engineer will implement a predeptoyment step that runs some security checks over the CloudFormation template before the pipeline processes it. This check should allow only inbound security group rules with a source CIDR of 0.0.0.0/0 if the rule has the description "Security Approval Ref XXXXX (where XXXXX is a preallocated reference). The pipeline step should fail if this condition is not met and the deployment should be blocked
How should this be accomplished?

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

Answer: B

NEW QUESTION 116
You have decided that you need to change the instance type of your production instances which are running as part of an AutoScaling group. The entire architecture is deployed using CloudFormation Template. You currently have 4 instances in Production. You cannot have any interruption in service and need to ensure 2 instances are always runningduring the update? Which of the options below listed can be used for this?

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

Answer: A

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

NEW QUESTION 119
A DevOps engineer is currently running a container-based workload on-premises The engineer wants to move the application to AWS, but needs to keep the on-premises solution active because not all APIs will move at the same time. The traffic between AWS and the on-premises network should be secure and encrypted at all times. Low management overload is also a requirement.
Which combination of actions will meet these criteria? (Select THREE.)

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

Answer: BDF

NEW QUESTION 123
You have an ELB setup in AWS with EC2 instances running behind it. You have been requested to monitor the incoming connections to the ELB. Which of the below options can suffice this requirement?

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

Answer: B

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

NEW QUESTION 126
A DevOps Engineer uses Docker container technology to build an image-analysis application. The application often sees spikes in traffic. The Engineer must automatically scale the application in response to customer demand while maintaining cost effectiveness and minimizing any impact on availability.
What will allow the FASTEST response to spikes in traffic while fulfilling the other requirements?

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

G. Configure the ECS service to use Service Auto Scalin
H. Set up Amazon CloudWatch alarms to scale the ECS service and cluster.
I. Deploy containers on Amazon EC2 instance
J. Deploy a container scheduler to schedule containers onto EC2 instance
K. Configure EC2 Auto Scaling for EC2 instances based on available Amazon CloudWatch metrics.

Answer: A

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

NEW QUESTION 131
A development team manages website deployments using AWS CodeDeploy blue/green deployments. The application is running on Amazon EC2 instances behind an Application Load Balancer in an Auto Scaling group.
When deploying a new revision, the team notices the deployment eventually fails, but it takes a long time to fail. After further inspection, the team discovers the AllowTrafflc lifecycle event ran for an hour and eventually failed without providing any other information. The team wants to ensure failure notices are delivered more quickly while maintaining application availability even upon failure.
Which combination of actions should be taken to meet these requirements? (Select TWO.)

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

Answer: AE

NEW QUESTION 132
A DevOps Engineer just joined a new company that is already running workloads on Amazon EC2 instances. AWS has been adopted incrementally with no central governance. The Engineer must now assess how well the existing deployments comply with the following requirements:
*EC2 instances are running only approved AMIs.
*Amazon EBS volumes are encrypted.
*EC2 instances have an Owner tag.
*Root login over SSH is disabled on EC2 instances.
Which services should the Engineer use to perform this assessment with the LEAST amount of effort? (Select TWO.)

A. AWS Config
B. Amazon GuardDuty
C. AWS System Manager
D. AWS Directory Service
E. Amazon Inspector

Answer: AE

Explanation:
https://docs.aws.amazon.com/ja_jp/inspector/latest/userguide/inspector_security-best-practices.html

NEW QUESTION 137
An ecommerce company is running an application on AWS. The company wants to create a standby disaster recovery solution in an additional Region that keeps the current application code. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The database layer is hosted on an Amazon RDS MySQL Multi-AZ DB instance. Amazon Route 53 DNS records point to the ALB.
Which combination of actions will meet these requirements with the LOWEST cost? (Select THREE.)

A. Configure a failover routing policy for the application DNS entry.
B. Configure a geolocation routing policy for the application DNS entry.
C. Create a cross-Region RDS read replica in the new standby Region.
D. Migrate the database layer to Amazon DynamoDB and enable global replication to the new standby Region.
E. Provision the ALB and Auto Scaling group in the new standby Region and set the desired capacity to match the active Region.
F. Provision the ALB and Auto Scaling group in the new standby Region and set the desired capacity to 1.

Answer: AEF

NEW QUESTION 139
Am 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 142
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 Configuratio
B. Create a new AMI with the updated ap
C. Use Route53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs.
D. Create new AM Is with the new ap
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 Autoscaling 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 Clastic beanstalk is good for development environments, and also there is no mention of having 2 environments where environment url's
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 146
You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty. What is the reason for this.

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

Answer: D

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

NEW QUESTION 151
A DevOps Engineer has several legacy applications that all generate different log formats. The Engineer must standardize the formats before writing them to Amazon S3 for querying and analysis. How can this requirement be met at the LOWEST cost?

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

Answer: D

NEW QUESTION 153
A web application has been deployed using an AWS Elastic Beanstalk application The Application Developers are concerned that they are seeing high latency in two different areas of the application: HTTP client requests to a third-party API MySQL client library queries to an Amazon RDS database A DevOps Engineer must gather trace data to diagnose the issues. Which steps will gather the trace information with the LEAST amount of changes and performance impacts to the application?

A. Add additional logging to the application cod
B. Use the Amazon CloudWatch agent to stream the application logs into Amazon Elasticsearch Servic
C. Query the log data in Amazon ES.
D. Instrument the application to use the AWS X-Ray SD
E. Post trace data to an Amazon Elasticsearch Service cluste
F. Query the trace data for calls to the HTTP client and the MySQL client.

G. On the AWS Elastic Beanstalk management page for the application, enable the AWS X-Ray daemon.View the trace data in the X-Ray console.
H. Instrument the application using the AWS X-Ray SD
I. On the AWS Elastic Beanstalk management page for the application, enable the X-Ray daemo
J. View the trace data in the X-Ray console.

Answer: C

NEW QUESTION 156
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 templat
B. Use the diversified allocation strategy with step scaling behind the ALB.
C. Add a Spot block resource to the AWS CloudFormation templat
D. Use the lowest-price allocation strategy with target tracking scaling behind the ALB.
E. Add a Spot Fleet resource to the AWS CloudFormation templat
F. Use the capacity-optimized allocation strategy with step scaling behind the ALB.
G. Add a Spot Fleet resource to the AWS CloudFormation templat
H. Use the diversified allocation strategy with scheduled scaling behind the ALB

Answer: C

NEW QUESTION 157
A DevOps Engineer needs to deploy a scalable three-tier Node.js application in AWS. The application must have zero downtime during deployments and be able to roll back to previous versions. Other applications will also connect to the same MySQL backend database.
The CIO has provided the following guidance for logging:
*Centrally view all current web access server logs.
*Search and filter web and application logs in near-real time.
*Retain log data for three months.
How should these requirements be met?

A. Deploy the application using AWS Elastic Beanstal
B. Configure the environment type for Elastic Load Balancing and Auto Scalin
C. Create an Amazon RDS MySQL instance inside the Elastic Beanstalk stac
D. Configure the Elastic Beanstalk log options to stream logs to Amazon CloudWatch Log
E. Set retention to 90 days.
F. Deploy the application on Amazon EC2. Configure Elastic Load Balancing and Auto Scalin
G. Use an Amazon RDS MySQL instance for the database tie
H. Configure the application to store log files in Amazon S3. Use Amazon EMR to search and filter the dat
I. Set an Amazon S3 lifecycle rule to expire objects after 90 days.
J. Deploy the application using AWS Elastic Beanstal
K. Configure the environment type for Elastic Load Balancing and Auto Scalin
L. Create the Amazon RDS MySQL instance outside the Elastic Beanstalk stac
M. Configure the Elastic Beanstalk log options to stream logs to Amazon CloudWatch Log
N. Set retention to 90 days.
O. Deploy the application on Amazon EC2. Configure Elastic Load Balancing and Auto Scalin
P. Use an Amazon RDS MySQL instance for the database tie
Q. Configure the application to load streaming log data using Amazon Kinesis Data Firehouse into Amazon E
R. Delete and create a new Amazon ES domain every 90 days.

Answer: B

Explanation:
https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-debugging.html

NEW QUESTION 158
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 160
A DevOps engineer is setting up a container-based architecture. The engineer has decided to use AWS CloudFormation to automatically provision an Amazon ECS cluster and an Amazon EC2 Auto Scaling group to launch the EC2 container instances. After successfully creating the CloudFormation stack, the engineer noticed that, even though the ECS cluster and the EC2 instances were created successfully and the stack finished the creation, the EC2 instances were associating with a different cluster.
How should the DevOps engineer update the CloudFormation template to resolve this issue?

A. Reference the EC2 instances in the AWS::ECS::Cluster resource and reference the ECS cluster in the AWS::ECS::Service resource.
B. Reference the ECS cluster in the AWS::AutoScaling::LaunchConfiguration resource of the UserData property.
C. Reference the ECS cluster in the AWS::EC2::Instance resource of the UserData property.
D. Reference the ECS cluster in the AWS::CloudFormation::CustomResource resource to trigger an AWS Lambda function that registers the EC2 instances with the appropriate ECS cluster.

Answer: B

NEW QUESTION 164
A DevOps Engineer is responsible for the deployment of a PHP application. The Engineer is working in a hybrid deployment, with the application running on both on-premises servers and Amazon EC2 instances. The application needs access to a database containing highly confidential information. Application instances need access to database credentials, which must be encrypted at rest and in transit before reaching the instances.
How should the Engineer automate the deployment process while also meeting the security requirements?

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

Answer: B

NEW QUESTION 168
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

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