



Amazon-Web-Services

Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional

About ExamBible

Your Partner of IT Exam

Found in 1998

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

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

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

NEW QUESTION 1

A company uses a single AWS account to test applications on Amazon EC2 instances. The company has turned on AWS Config in the AWS account and has activated the restricted-ssh AWS Config managed rule.

The company needs an automated monitoring solution that will provide a customized notification in real time if any security group in the account is not compliant with the restricted-ssh rule. The customized notification must contain the name and ID of the noncompliant security group.

A DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic in the account and subscribes the appropriate personnel to the topic. What should the DevOps engineer do next to meet these requirements?

- A. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule
- B. Configure an input transformer for the EventBridge rule
- C. Configure the EventBridge rule to publish a notification to the SNS topic
- D. Configure AWS Config to send all evaluation results for the restricted-ssh rule to the SNS topic
- E. Configure a filter policy on the SNS topic to send only notifications that contain the text of NON_COMPLIANT in the notification to subscribers.
- F. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule
- G. Configure the EventBridge rule to invoke AWS Systems Manager Run Command on the SNS topic to customize a notification and to publish the notification to the SNS topic
- H. Create an Amazon EventBridge rule that matches all AWS Config evaluation results of NON_COMPLIANT
- I. Configure an input transformer for the restricted-ssh rule
- J. Configure the EventBridge rule to publish a notification to the SNS topic

Answer: A

Explanation:

Create an Amazon EventBridge (Amazon CloudWatch Events) rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure an input transformer for the EventBridge (CloudWatch Events) rule. Configure the EventBridge (CloudWatch Events) rule to publish a notification to the SNS topic. This approach uses Amazon EventBridge (previously known as Amazon CloudWatch Events) to filter AWS Config evaluation results based on the restricted-ssh rule and its compliance status (NON_COMPLIANT). An input transformer can be used to customize the information contained in the notification, such as the name and ID of the noncompliant security group. The EventBridge (CloudWatch Events) rule can then be configured to publish a notification to the SNS topic, which will notify the appropriate personnel in real-time.

NEW QUESTION 2

A company is examining its disaster recovery capability and wants the ability to switch over its daily operations to a secondary AWS Region. The company uses AWS CodeCommit as a source control tool in the primary Region.

A DevOps engineer must provide the capability for the company to develop code in the secondary Region. If the company needs to use the secondary Region, developers can add an additional remote URL to their local Git configuration.

Which solution will meet these requirements?

- A. Create a CodeCommit repository in the secondary Region
- B. Create an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's CodeCommit repository
- C. Create an AWS Lambda function that invokes the CodeBuild project
- D. Create an Amazon EventBridge rule that reacts to merge events in the primary Region's CodeCommit repository
- E. Configure the EventBridge rule to invoke the Lambda function.
- F. Create an Amazon S3 bucket in the secondary Region
- G. Create an AWS Fargate task to perform a Git mirror operation of the primary Region's CodeCommit repository and copy the result to the S3 bucket
- H. Create an AWS Lambda function that initiates the Fargate task
- I. Create an Amazon EventBridge rule that reacts to merge events in the CodeCommit repository
- J. Configure the EventBridge rule to invoke the Lambda function.
- K. Create an AWS CodeArtifact repository in the secondary Region
- L. Create an AWS CodePipeline pipeline that uses the primary Region's CodeCommit repository for the source action
- M. Create a Cross-Region stage in the pipeline that packages the CodeCommit repository contents and stores the contents in the CodeArtifact repository when a pull request is merged into the CodeCommit repository.
- N. Create an AWS Cloud9 environment and a CodeCommit repository in the secondary Region
- O. Configure the primary Region's CodeCommit repository as a remote repository in the AWS Cloud9 environment
- P. Connect the secondary Region's CodeCommit repository to the AWS Cloud9 environment.

Answer: A

Explanation:

The best solution to meet the disaster recovery capability and allow developers to switch over to a secondary AWS Region for code development is option A. This involves creating a CodeCommit repository in the secondary Region and setting up an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's repository. An AWS Lambda function is then created to invoke the CodeBuild project. Additionally, an Amazon EventBridge rule is configured to react to merge events in the primary Region's CodeCommit repository and invoke the Lambda function. This setup ensures that the secondary Region's repository is always up-to-date with the primary repository, allowing for a seamless transition in case of a disaster recovery event.

References:

? AWS CodeCommit User Guide on resilience and disaster recovery.

? AWS Documentation on monitoring CodeCommit events in Amazon EventBridge and Amazon CloudWatch Events.

NEW QUESTION 3

A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.

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

- A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store. Use Aurora's automatic recovery capabilities in the event of a disaster.
- B. Create an Amazon Aurora global database in two Regions as the data store.
- C. In the event of a failure, promote the secondary Region as the primary for the application.
- D. Create an Amazon Aurora multi-master cluster across multiple Regions as the data store.
- E. Use a Network Load Balancer to balance the database traffic in different Regions.
- F. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Regions.
- G. Use health checks to determine the availability in a given Region.

- H. Use Auto Scaling groups in each Region to adjust capacity based on demand.
- I. Set up the application in two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on demand.
- J. In the event of a disaster adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

Answer: BD

NEW QUESTION 4

A DevOps engineer is implementing governance controls for a company that requires its infrastructure to be housed within the United States. The engineer must restrict which AWS Regions can be used, and ensure an alert is sent as soon as possible if any activity outside the governance policy takes place. The controls should be automatically enabled on any new Region outside the United States (US).

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

- A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Region
- B. Attach the policy to the root of the organization.
- C. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Region
- D. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions.
- E. Use an AWS Lambda function that checks for AWS service activity and deploy it to all Region
- F. Write an Amazon EventBridge rule that runs the Lambda function every hour, sending an alert if activity is found in a non-US Region.
- G. Use an AWS Lambda function to query Amazon Inspector to look for service activity in non-US Regions and send alerts if any activity is found.
- H. Write an SCP using the `aws:RequestedRegion` condition key limiting access to US Region
- I. Apply the policy to all users, groups, and roles

Answer: AB

Explanation:

To implement governance controls that restrict AWS service usage to within the United States and ensure alerts for any activity outside the governance policy, the following actions will meet the requirements:

? A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Regions. Attach the policy to the root of the organization. This action will effectively prevent users and roles in all accounts within the organization from accessing services in non-US Regions¹².

? B. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Regions. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions. This action will allow monitoring of all AWS Regions and will trigger alerts if any activity is detected in non-US Regions, ensuring that the governance team is notified as soon as possible³.

References:

? AWS Documentation on Service Control Policies (SCPs) and how they can be used to manage permissions and restrict access based on Regions¹².

? AWS Documentation on monitoring CloudTrail log files with Amazon CloudWatch Logs to set up alerts for specific activities³.

NEW QUESTION 5

A company has an application that runs on AWS Lambda and sends logs to Amazon CloudWatch Logs. An Amazon Kinesis data stream is subscribed to the log groups in CloudWatch Logs. A single consumer Lambda function processes the logs from the data stream and stores the logs in an Amazon S3 bucket.

The company's DevOps team has noticed high latency during the processing and ingestion of some logs.

Which combination of steps will reduce the latency? (Select THREE.)

- A. Create a data stream consumer with enhanced fan-out
- B. Set the Lambda function that processes the logs as the consumer.
- C. Increase the `ParallelizationFactor` setting in the Lambda event source mapping.
- D. Configure reserved concurrency for the Lambda function that processes the logs.
- E. Increase the batch size in the Kinesis data stream.
- F. Turn off the `ReportBatchItemFailures` setting in the Lambda event source mapping.
- G. Increase the number of shards in the Kinesis data stream.

Answer: ABC

Explanation:

The latency in processing and ingesting logs can be caused by several factors, such as the throughput of the Kinesis data stream, the concurrency of the Lambda function, and the configuration of the event source mapping. To reduce the latency, the following steps can be taken:

? Create a data stream consumer with enhanced fan-out. Set the Lambda function that processes the logs as the consumer. This will allow the Lambda function to receive records from the data stream with dedicated throughput of up to 2 MB per second per shard, independent of other consumers¹. This will reduce the contention and delay in accessing the data stream.

? Increase the `ParallelizationFactor` setting in the Lambda event source mapping. This will allow the Lambda service to invoke more instances of the function concurrently to process the records from the data stream². This will increase the processing capacity and reduce the backlog of records in the data stream.

? Configure reserved concurrency for the Lambda function that processes the logs. This will ensure that the function has enough concurrency available to handle the increased load from the data stream³. This will prevent the function from being throttled by the account-level concurrency limit.

The other options are not effective or may have negative impacts on the latency. Option D is not suitable because increasing the batch size in the Kinesis data stream will increase the amount of data that the Lambda function has to process in each invocation, which may increase the execution time and latency⁴. Option E is not advisable because turning off the `ReportBatchItemFailures` setting in the Lambda event source mapping will prevent the Lambda service from retrying the failed records, which may result in data loss. Option F is not necessary because increasing the number of shards in the Kinesis data stream will increase the throughput of the data stream, but it will not affect the processing speed of the Lambda function, which is the bottleneck in this scenario.

References:

? 1: Using AWS Lambda with Amazon Kinesis Data Streams - AWS Lambda

? 2: AWS Lambda event source mappings - AWS Lambda

? 3: Managing concurrency for a Lambda function - AWS Lambda

? 4: AWS Lambda function scaling - AWS Lambda

? : AWS Lambda event source mappings - AWS Lambda

? : Scaling Amazon Kinesis Data Streams with AWS CloudFormation - Amazon Kinesis Data Streams

NEW QUESTION 6

A company deploys a web application on Amazon EC2 instances that are behind an

Application Load Balancer (ALB). The company stores the application code in an AWS CodeCommit repository. When code is merged to the main branch, an AWS Lambda function invokes an AWS CodeBuild project. The CodeBuild project packages the code, stores the packaged code in AWS CodeArtifact, and invokes AWS Systems Manager Run Command to deploy the packaged code to the EC2 instances.

Previous deployments have resulted in defects, EC2 instances that are not running the latest version of the packaged code, and inconsistencies between instances.

Which combination of actions should a DevOps engineer take to implement a more reliable deployment solution? (Select TWO.)

- A. Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provide
- B. Configure pipeline stages that run the CodeBuild project in parallel to build and test the applicatio
- C. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.
- D. Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provide
- E. Create separate pipeline stages that run a CodeBuild project to build and then test the applicatio
- F. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.
- G. Create an AWS CodeDeploy application and a deployment group to deploy the packaged code to the EC2 instance
- H. Configure the ALB for the deployment group.
- I. Create individual Lambda functions that use AWS CodeDeploy instead of Systems Manager to run build, test, and deploy actions.
- J. Create an Amazon S3 bucket
- K. Modify the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifact
- L. Use deploy actions in CodeDeploy to deploy the artifact to the EC2 instances.

Answer: AC

Explanation:

To implement a more reliable deployment solution, a DevOps engineer should take the following actions:

? Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provider. Configure pipeline stages that run the CodeBuild project in parallel to build and test the application. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action. This action will improve the deployment reliability by automating the entire process from code commit to deployment, reducing human errors and inconsistencies. By running the build and test stages in parallel, the pipeline can also speed up the delivery time and provide faster feedback. By using CodeDeploy as the deployment action, the pipeline can leverage the features of CodeDeploy, such as traffic shifting, health checks, rollback, and deployment configuration¹²³

? Create an AWS CodeDeploy application and a deployment group to deploy the packaged code to the EC2 instances. Configure the ALB for the deployment group. This action will improve the deployment reliability by using CodeDeploy to orchestrate the deployment across multiple EC2 instances behind an ALB. CodeDeploy can perform blue/green deployments or in-place deployments with traffic shifting, which can minimize downtime and reduce risks. CodeDeploy can also monitor the health of the instances during and after the deployment, and automatically roll back if any issues are detected. By configuring the ALB for the deployment group, CodeDeploy can register and deregister instances from the load balancer as needed, ensuring that only healthy instances receive traffic⁴⁵

The other options are not correct because they do not improve the deployment reliability or follow best practices. Creating separate pipeline stages that run a CodeBuild project to build and then test the application is not a good option because it will increase the pipeline execution time and delay the feedback loop. Creating individual Lambda functions that use CodeDeploy instead of Systems Manager to run build, test, and deploy actions is not a valid option because it will add unnecessary complexity and cost to the solution. Lambda functions are not designed for long-running tasks such as building or deploying applications. Creating an Amazon S3 bucket and modifying the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifact is not a necessary option because it will not affect the deployment reliability. CodeArtifact is a secure, scalable, and cost-effective package management service that can store and share software packages for application development⁶⁷

References:

- ? 1: What is AWS CodePipeline? - AWS CodePipeline
- ? 2: Create a pipeline in AWS CodePipeline - AWS CodePipeline
- ? 3: Deploy an application with AWS CodeDeploy - AWS CodePipeline
- ? 4: What is AWS CodeDeploy? - AWS CodeDeploy
- ? 5: Configure an Application Load Balancer for your blue/green deployments - AWS CodeDeploy
- ? 6: What is AWS Lambda? - AWS Lambda
- ? 7: What is AWS CodeArtifact? - AWS CodeArtifact

NEW QUESTION 7

A DevOps engineer is building an application that uses an AWS Lambda function to query an Amazon Aurora MySQL DB cluster. The Lambda function performs only read queries. Amazon EventBridge events invoke the Lambda function.

As more events invoke the Lambda function each second, the database's latency increases and the database's throughput decreases. The DevOps engineer needs to improve the performance of the application.

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

- A. Use Amazon RDS Proxy to create a proxy
- B. Connect the proxy to the Aurora cluster reader endpoint
- C. Set a maximum connections percentage on the proxy.
- D. Implement database connection pooling inside the Lambda code
- E. Set a maximum number of connections on the database connection pool.
- F. Implement the database connection opening outside the Lambda event handler code.
- G. Implement the database connection opening and closing inside the Lambda event handler code.
- H. Connect to the proxy endpoint from the Lambda function.
- I. Connect to the Aurora cluster endpoint from the Lambda function.

Answer: ACE

Explanation:

To improve the performance of the application, the DevOps engineer should use Amazon RDS Proxy, implement the database connection opening outside the Lambda event handler code, and connect to the proxy endpoint from the Lambda function. References:

? Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure¹. By using Amazon RDS Proxy, the DevOps engineer can reduce the overhead of opening and closing connections to the database, which can improve latency and throughput².

? The DevOps engineer should connect the proxy to the Aurora cluster reader endpoint, which allows read-only connections to one of the Aurora Replicas in the DB cluster³. This can help balance the load across multiple read replicas and improve performance for read-intensive workloads⁴.

? The DevOps engineer should implement the database connection opening outside the Lambda event handler code, which means using a global variable to store the database connection object⁵. This can enable connection reuse across multiple invocations of the Lambda function, which can reduce latency and improve performance.

? The DevOps engineer should connect to the proxy endpoint from the Lambda function, which is a unique URL that represents the proxy. This can allow the Lambda function to access the database through the proxy, which can provide benefits such as connection pooling, load balancing, failover handling, and enhanced security.

? The other options are incorrect because:

NEW QUESTION 8

A company uses AWS Key Management Service (AWS KMS) keys and manual key rotation to meet regulatory compliance requirements. The security team wants to be notified when any keys have not been rotated after 90 days.

Which solution will accomplish this?

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

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-aws-config-to-determine-compliance-of-aws-kms-key-policies-to-your-specifications/>

NEW QUESTION 9

A development team is using AWS CodeCommit to version control application code and AWS CodePipeline to orchestrate software deployments. The team has decided to use a remote main branch as the trigger for the pipeline to integrate code changes. A developer has pushed code changes to the CodeCommit repository, but noticed that the pipeline had no reaction, even after 10 minutes.

Which of the following actions should be taken to troubleshoot this issue?

- A. Check that an Amazon EventBridge rule has been created for the main branch to trigger the pipeline.
- B. Check that the CodePipeline service role has permission to access the CodeCommit repository.
- C. Check that the developer's IAM role has permission to push to the CodeCommit repository.
- D. Check to see if the pipeline failed to start because of CodeCommit errors in Amazon CloudWatch Logs.

Answer: A

Explanation:

When you create a pipeline from CodePipeline during the step-by-step it creates a CloudWatch Event rule for a given branch and repo like this:

```
{
  "source": [ "aws.codecommit"
],
  "detail-type": [
    "CodeCommit Repository State Change"
  ],
  "resources": [
    "arn:aws:codecommit:us-east-1:xxxxx:repo-name"
  ],
  "detail": {
    "event": [ "referenceCreated", "referenceUpdated"
  ],
  "referenceType": [ "branch"
  ],
  "referenceName": [ "master"
  ]
}
```

<https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-trigger-source-repo-changes-console.html>

NEW QUESTION 10

A company has a data ingestion application that runs across multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to monitor the application and consolidate access to the application. Currently the company is running the application on Amazon EC2 instances from several Auto Scaling groups. The EC2 instances have no access to the internet because the data is sensitive. Engineers have deployed the necessary VPC endpoints. The EC2 instances run a custom AMI that is built specifically for the application.

To maintain and troubleshoot the application, system administrators need the ability to log in to the EC2 instances. This access must be automated and controlled centrally. The company's security team must receive a notification whenever the instances are accessed.

Which solution will meet these requirements?

- A. Create an Amazon EventBridge rule to send notifications to the security team whenever a user logs in to an EC2 instance. Use EC2 Instance Connect to log in to the instance.
- B. Deploy Auto Scaling groups by using AWS CloudFormation. Use the cfn-init helper script to deploy appropriate VPC routes for external access. Rebuild the custom AMI so that the custom AMI includes AWS Systems Manager Agent.
- C. Deploy a NAT gateway and a bastion host that has internet access. Create a security group that allows incoming traffic on all the EC2 instances from the bastion host. Install AWS Systems Manager Agent on all the EC2 instances. Use Auto Scaling group lifecycle hooks for monitoring and auditing access. Use Systems Manager Session Manager to log into the instances. Send logs to a log group in Amazon CloudWatch Log.
- D. Export data to Amazon S3 for auditing. Send notifications to the security team by using S3 event notifications.
- E. Use EC2 Image Builder to rebuild the custom AMI. Include the most recent version of AWS Systems Manager Agent in the image. Configure the Auto Scaling group to attach the AmazonSSMManagedInstanceCore role to all the EC2 instances. Use Systems Manager Session Manager to log in to the instances. Enable logging of session details to Amazon S3. Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.
- F. Use AWS Systems Manager Automation to build Systems Manager Agent into the custom AMI. Configure AWS Config to attach an SCP to the root organization account to allow the EC2 instances to connect to Systems Manager. Use Systems Manager Session Manager to log in to the instances. Enable logging of session details to Amazon S3. Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: C

Explanation:

Even if AmazonSSMManagedInstanceCore is a managed policy and not an IAM role I will go with C because this policy is to be attached to an IAM role for EC2 to access System Manager.

NEW QUESTION 10

A company has a new AWS account that teams will use to deploy various applications. The teams will create many Amazon S3 buckets for application- specific purposes and to store AWS CloudTrail logs. The company has enabled Amazon Macie for the account.

A DevOps engineer needs to optimize the Macie costs for the account without compromising the account's functionality.

Which solutions will meet these requirements? (Select TWO.)

- A. Exclude S3 buckets that contain CloudTrail logs from automated discovery.
- B. Exclude S3 buckets that have public read access from automated discovery.
- C. Configure scheduled daily discovery jobs for all S3 buckets in the account.
- D. Configure discovery jobs to include S3 objects based on the last modified criterion.
- E. Configure discovery jobs to include S3 objects that are tagged as production only.

Answer: AD

Explanation:

To optimize the Macie costs for the account without compromising the account's functionality, the DevOps engineer needs to exclude S3 buckets that do not contain sensitive data from automated discovery. S3 buckets that contain CloudTrail logs are unlikely to have sensitive data, and Macie charges for scanning and monitoring data in S3 buckets. Therefore, excluding S3 buckets that contain CloudTrail logs from automated discovery can reduce Macie costs. Similarly, configuring discovery jobs to include S3 objects based on the last modified criterion can also reduce Macie costs, as it will only scan and monitor new or updated objects, rather than all objects in the bucket.

NEW QUESTION 14

A company has a mobile application that makes HTTP API calls to an Application Load Balancer (ALB). The ALB routes requests to an AWS Lambda function. Many different versions of the application are in use at any given time, including versions that are in testing by a subset of users. The version of the application is defined in the user-agent header that is sent with all requests to the API.

After a series of recent changes to the API, the company has observed issues with the application. The company needs to gather a metric for each API operation by response code for each version of the application that is in use. A DevOps engineer has modified the Lambda function to extract the API operation name, version information from the user- agent header and response code.

Which additional set of actions should the DevOps engineer take to gather the required metrics?

- A. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group
- B. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name
- C. Specify response code and application version as dimensions for the metric.
- D. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group
- E. Configure a CloudWatch Logs Insights query to populate CloudWatch metrics from the log line
- F. Specify response code and application version as dimensions for the metric.
- G. Configure the ALB access logs to write to an Amazon CloudWatch Logs log group
- H. Modify the Lambda function to respond to the ALB with the API operation name, response code, and version number as response metadata
- I. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name
- J. Specify response code and application version as dimensions for the metric.
- K. Configure AWS X-Ray integration on the Lambda function
- L. Modify the Lambda function to create an X-Ray subsegment with the API operation name, response code, and version number
- M. Configure X-Ray insights to extract an aggregated metric for each API operation name and to publish the metric to Amazon CloudWatch
- N. Specify response code and application version as dimensions for the metric.

Answer: A

Explanation:

"Note that the metric filter is different from a log insights query, where the experience is interactive and provides immediate search results for the user to investigate.

No automatic action can be invoked from an insights query. Metric filters, on the other hand, will generate metric data in the form of a time series. This lets you create alarms that integrate into your ITSM processes, execute AWS Lambda functions, or even create anomaly detection models."

<https://aws.amazon.com/blogs/mt/quantify-custom-application- metrics-with-amazon-cloudwatch-logs-and-metric-filters/>

NEW QUESTION 19

A production account has a requirement that any Amazon EC2 instance that has been logged in to manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with the Amazon CloudWatch Logs agent configured.

How can this process be automated?

- A. Create a CloudWatch Logs subscription to an AWS Step Functions application
- B. Configure an AWS Lambda function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned
- C. Create an Amazon EventBridge rule to invoke a second Lambda function once a day that will terminate all instances with this tag.
- D. Create an Amazon CloudWatch alarm that will be invoked by the login event
- E. Send the notification to an Amazon Simple Notification Service (Amazon SNS) topic that the operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.
- F. Create an Amazon CloudWatch alarm that will be invoked by the login event
- G. Configure the alarm to send to an Amazon Simple Queue Service (Amazon SQS) queue
- H. Use a group of worker instances to process messages from the queue, which then schedules an Amazon EventBridge rule to be invoked.
- I. Create a CloudWatch Logs subscription to an AWS Lambda function
- J. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned
- K. Create an Amazon EventBridge rule to invoke a daily Lambda function that terminates all instances with this tag.

Answer: D

Explanation:

"You can use subscriptions to get access to a real-time feed of log events from CloudWatch Logs and have it delivered to other services such as an Amazon

Kinesis stream, an Amazon Kinesis Data Firehose stream, or AWS Lambda for custom processing, analysis, or loading to other systems. When log events are sent to the receiving service, they are Base64 encoded and compressed with the gzip format." See <https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/Subscriptions.html>

NEW QUESTION 20

A company recently launched multiple applications that use Application Load Balancers. Application response time often slows down when the applications experience problems. A DevOps engineer needs to implement a monitoring solution that alerts the company when the applications begin to perform slowly. The DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic and subscribes the company's email address to the topic. What should the DevOps engineer do next to meet the requirements?

- A. Create an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval. Configure the Lambda function to publish a notification to the SNS topic when the applications return errors.
- B. Create an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interval.
- C. Configure the canary to use the SNS topic when the applications return errors.
- D. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric. Configure the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports. Configure the CloudWatch alarm to use the SNS topic.
- E. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric. Configure the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports. Configure the CloudWatch alarm to use the SNS topic.

Answer: B

Explanation:

? Option A is incorrect because creating an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval is not a valid solution. EventBridge rules can only trigger Lambda functions based on events, not on time intervals. Moreover, querying the applications on a 5-minute interval might incur unnecessary costs and network overhead, and might not detect performance issues in real time.

? Option B is correct because creating an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interval is a valid solution. CloudWatch Synthetics canaries are configurable scripts that monitor endpoints and APIs by simulating customer behavior. Canaries can run as often as once per minute, and can measure the latency and availability of the applications. Canaries can also send notifications to an Amazon SNS topic when they detect errors or performance issues¹.

? Option C is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution. The RequestCountPerTarget metric measures the number of requests completed or connections made per target in a target group². This metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports is not a valid way to measure the application performance, as it depends on the application design and implementation.

? Option D is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution, for the same reason as option C. The RequestCountPerTarget metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports is not a valid way to measure the application performance, as it does not account for variability or outliers in the response time distribution.

References:

? 1: Using synthetic monitoring

? 2: Application Load Balancer metrics

NEW QUESTION 24

A company manages multiple AWS accounts by using AWS Organizations with OUs for the different business divisions. The company is updating their corporate network to use new IP address ranges. The company has 10 Amazon S3 buckets in different AWS accounts. The S3 buckets store reports for the different divisions. The S3 bucket configurations allow only private corporate network IP addresses to access the S3 buckets.

A DevOps engineer needs to change the range of IP addresses that have permission to access the contents of the S3 buckets. The DevOps engineer also needs to revoke the permissions of two OUs in the company.

Which solution will meet these requirements?

- A. Create a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that denies access to the old range of IP addresses for all the S3 bucket.
- B. Set a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets.
- C. Create a new SCP that has a statement that allows only the new range of IP addresses to access the S3 bucket.
- D. Create another SCP that denies access to the S3 bucket.
- E. Attach the second SCP to the two OUs.
- F. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket.
- G. Create a new SCP that denies access to the S3 bucket.
- H. Attach the SCP to the two OUs.
- I. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket.
- J. Set a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets.

Answer: C

Explanation:

The correct answer is C.

A comprehensive and detailed explanation is:

? Option A is incorrect because creating a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that denies access to the old range of IP addresses for all the S3 buckets, is not a valid solution. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. SCPs can only control the actions that can be performed by the principals in the organization, not the access to specific resources. Moreover, setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets.

? Option B is incorrect because creating a new SCP that has a statement that allows

only the new range of IP addresses to access the S3 buckets is not a valid solution, for the same reason as option A. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. Creating another SCP that denies access to the S3 buckets and attaching it to the two OUs is also not a valid solution, as SCPs cannot specify the S3 buckets as resources either.

? Option C is correct because it meets both requirements of changing the range of IP addresses that have permission to access the contents of the S3 buckets and revoking the permissions of two OUs in the company. On all the S3 buckets, configuring resource-based policies that allow only the new range of IP

addresses to access the S3 buckets is a valid way to update the IP address ranges, as resource-based policies can specify both resources and conditions. Creating a new SCP that denies access to the S3 buckets and attaching it to the two OUs is also a valid way to revoke the permissions of those OUs, as SCPs can deny actions such as s3:PutObject or s3:GetObject on any resource.

? Option D is incorrect because setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets. A permissions boundary is a policy that defines the maximum permissions that an IAM entity can have. However, it does not revoke any existing permissions that are granted by other policies.

References:

- ? AWS Organizations
- ? S3 Bucket Policies
- ? Service Control Policies
- ? Permissions Boundaries

NEW QUESTION 27

A healthcare services company is concerned about the growing costs of software licensing for an application for monitoring patient wellness. The company wants to create an audit process to ensure that the application is running exclusively on Amazon EC2 Dedicated Hosts. A DevOps engineer must create a workflow to audit the application to ensure compliance.

What steps should the engineer take to meet this requirement with the LEAST administrative overhead?

- A. Use AWS Systems Manager Configuration Compliance
- B. Use calls to the put-compliance-items API action to scan and build a database of noncompliant EC2 instances based on their host placement configuration
- C. Use an Amazon DynamoDB table to store these instance IDs for fast access
- D. Generate a report through Systems Manager by calling the list-compliance-summaries API action.
- E. Use custom Java code running on an EC2 instance
- F. Set up EC2 Auto Scaling for the instance depending on the number of instances to be checked
- G. Send the list of noncompliant EC2 instance IDs to an Amazon SQS queue
- H. Set up another worker instance to process instance IDs from the SQS queue and write them to Amazon DynamoDB
- I. Use an AWS Lambda function to terminate noncompliant instance IDs obtained from the queue, and send them to an Amazon SNS email topic for distribution.
- J. Use AWS Config
- K. Identify all EC2 instances to be audited by enabling Config Recording on all Amazon EC2 resources for the region
- L. Create a custom AWS Config rule that triggers an AWS Lambda function by using the "config-rule-change-triggered" blueprint. Modify the Lambda evaluateCompliance () function to verify host placement to return a NON_COMPLIANT result if the instance is not running on an EC2 Dedicated Host
- M. Use the AWS Config report to address noncompliant instances.
- N. Use AWS CloudTrail
- O. Identify all EC2 instances to be audited by analyzing all calls to the EC2 RunCommand API action
- P. Invoke a AWS Lambda function that analyzes the host placement of the instance
- Q. Store the EC2 instance ID of noncompliant resources in an Amazon RDS for MySQL DB instance
- R. Generate a report by querying the RDS instance and exporting the query results to a CSV text file.

Answer: C

Explanation:

The correct answer is C. Using AWS Config to identify and audit all EC2 instances based on their host placement configuration is the most efficient and scalable solution to ensure compliance with the software licensing requirement. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. By creating a custom AWS Config rule that triggers a Lambda function to verify host placement, the DevOps engineer can automate the process of checking whether the instances are running on EC2 Dedicated Hosts or not. The Lambda function can return a NON_COMPLIANT result if the instance is not running on an EC2 Dedicated Host, and the AWS Config report can provide a summary of the compliance status of the instances. This solution requires the least administrative overhead compared to the other options.

Option A is incorrect because using AWS Systems Manager Configuration Compliance to scan and build a database of noncompliant EC2 instances based on their host placement configuration is a more complex and costly solution than using AWS Config. AWS Systems Manager Configuration Compliance is a feature of AWS Systems Manager that enables you to scan your managed instances for patch compliance and configuration inconsistencies. To use this feature, the DevOps engineer would need to install the Systems Manager Agent on each EC2 instance, create a State Manager association to run the put-compliance-items API action periodically, and use a DynamoDB table to store the instance IDs of noncompliant resources. This solution would also require more API calls and storage costs than using AWS Config.

Option B is incorrect because using custom Java code running on an EC2 instance to check and terminate noncompliant EC2 instances is a more cumbersome and error-prone solution than using AWS Config. This solution would require the DevOps engineer to write and maintain the Java code, set up EC2 Auto Scaling for the instance, use an SQS queue and another worker instance to process the instance IDs, use a Lambda function and an SNS topic to terminate and notify the noncompliant instances, and handle any potential failures or exceptions in the workflow. This solution would also incur more compute, storage, and messaging costs than using AWS Config.

Option D is incorrect because using AWS CloudTrail to identify and audit EC2 instances by analyzing the EC2 RunCommand API action is a less reliable and accurate solution than using AWS Config. AWS CloudTrail is a service that enables you to monitor and log the API activity in your AWS account. The EC2 RunCommand API action is used to execute commands on one or more EC2 instances. However, this API action does not necessarily indicate the host placement of the instance, and it may not capture all the instances that are running on EC2 Dedicated Hosts or not. Therefore, option D would not provide a comprehensive and consistent audit of the EC2 instances.

NEW QUESTION 30

A company's developers use Amazon EC2 instances as remote workstations. The company is concerned that users can create or modify EC2 security groups to allow unrestricted inbound access.

A DevOps engineer needs to develop a solution to detect when users create unrestricted security group rules. The solution must detect changes to security group rules in near real time, remove unrestricted rules, and send email notifications to the security team. The DevOps engineer has created an AWS Lambda function that checks for security group ID from input, removes rules that grant unrestricted access, and sends notifications through Amazon Simple Notification Service (Amazon SNS).

What should the DevOps engineer do next to meet the requirements?

- A. Configure the Lambda function to be invoked by the SNS topic
- B. Create an AWS CloudTrail subscription for the SNS topic
- C. Configure a subscription filter for security group modification events.
- D. Create an Amazon EventBridge scheduled rule to invoke the Lambda function
- E. Define a schedule pattern that runs the Lambda function every hour.
- F. Create an Amazon EventBridge event rule that has the default event bus as the source
- G. Define the rule's event pattern to match EC2 security group creation and modification event

- H. Configure the rule to invoke the Lambda function.
- I. Create an Amazon EventBridge custom event bus that subscribes to events from all AWS service
- J. Configure the Lambda function to be invoked by the custom event bus.

Answer: C

Explanation:

To meet the requirements, the DevOps engineer should create an Amazon EventBridge event rule that has the default event bus as the source. The rule's event pattern should match EC2 security group creation and modification events, and it should be configured to invoke the Lambda function. This solution will allow for near real-time detection of security group rule changes and will trigger the Lambda function to remove any unrestricted rules and send email notifications to the security team. <https://repost.aws/knowledge-center/monitor-security-group-changes-ec2>

NEW QUESTION 32

A company has multiple member accounts that are part of an organization in AWS Organizations. The security team needs to review every Amazon EC2 security group and their inbound and outbound rules. The security team wants to programmatically retrieve this information from the member accounts using an AWS Lambda function in the management account of the organization.

Which combination of access changes will meet these requirements? (Choose three.)

- A. Create a trust relationship that allows users in the member accounts to assume the management account IAM role.
- B. Create a trust relationship that allows users in the management account to assume the IAM roles of the member accounts.
- C. Create an IAM role in each member account that has access to the AmazonEC2ReadOnlyAccess managed policy.
- D. Create an IAM role in each member account to allow the sts:AssumeRole action against the management account IAM role's ARN.
- E. Create an IAM role in the management account that allows the sts:AssumeRole action against the member account IAM role's ARN.
- F. Create an IAM role in the management account that has access to the AmazonEC2ReadOnlyAccess managed policy.

Answer: BCE

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/lambda-function-assume-iam-role/> <https://kreuzwerker.de/post/aws-multi-account-setups-reloaded>

NEW QUESTION 36

A DevOps engineer at a company is supporting an AWS environment in which all users use AWS IAM Identity Center (AWS Single Sign-On). The company wants to immediately disable credentials of any new IAM user and wants the security team to receive a notification.

Which combination of steps should the DevOps engineer take to meet these requirements? (Choose three.)

- A. Create an Amazon EventBridge rule that reacts to an IAM CreateUser API call in AWS CloudTrail.
- B. Create an Amazon EventBridge rule that reacts to an IAM GetLoginProfile API call in AWS CloudTrail.
- C. Create an AWS Lambda function that is a target of the EventBridge rule.
- D. Configure the Lambda function to disable any access keys and delete the login profiles that are associated with the IAM user.
- E. Create an AWS Lambda function that is a target of the EventBridge rule.
- F. Configure the Lambda function to delete the login profiles that are associated with the IAM user.
- G. Create an Amazon Simple Notification Service (Amazon SNS) topic that is a target of the EventBridge rule.
- H. Subscribe the security team's group email address to the topic.
- I. Create an Amazon Simple Queue Service (Amazon SQS) queue that is a target of the Lambda function.
- J. Subscribe the security team's group email address to the queue.

Answer: ACE

NEW QUESTION 40

A DevOps engineer has developed an AWS Lambda function. The Lambda function starts an AWS CloudFormation drift detection operation on all supported resources for a specific CloudFormation stack. The Lambda function then exits its invocation. The DevOps engineer has created an Amazon EventBridge scheduled rule that invokes the Lambda function every hour. An Amazon Simple Notification Service (Amazon SNS) topic already exists in the AWS account. The DevOps engineer has subscribed to the SNS topic to receive notifications.

The DevOps engineer needs to receive a notification as soon as possible when drift is detected in this specific stack configuration.

Which solution will meet these requirements?

- A. Configure the existing EventBridge rule to also target the SNS topic. Configure an SNS subscription filter policy to match the CloudFormation stack.
- B. Attach the subscription filter policy to the SNS topic.
- C. Create a second Lambda function to query the CloudFormation API for the drift detection results for the stack. Configure the second Lambda function to publish a message to the SNS topic. If drift is detected, adjust the existing EventBridge rule to also target the second Lambda function.
- D. Configure Amazon GuardDuty in the account with drift detection for all CloudFormation stacks.
- E. Create a second EventBridge rule that reacts to the GuardDuty drift detection event finding for the specific CloudFormation stack.
- F. Configure the SNS topic as a target of the second EventBridge rule.
- G. Configure AWS Config in the account.
- H. Use the cloudformation-stack-drift-detection-check managed rule.
- I. Create a second EventBridge rule that reacts to a compliance change event for the CloudFormation stack.
- J. Configure the SNS topic as a target of the second EventBridge rule.

Answer: D

Explanation:

A comprehensive and detailed explanation is:

? Option A is incorrect because EventBridge rules cannot filter events based on the message body or attributes of the target service. Therefore, configuring an SNS subscription filter policy to match the CloudFormation stack will not work. The SNS topic will receive all events from the EventBridge rule, regardless of the stack name or drift status.

? Option B is incorrect because it introduces unnecessary complexity and cost.

Creating a second Lambda function to query the CloudFormation API for the drift detection results is redundant, since CloudFormation already publishes drift detection events to EventBridge. Moreover, invoking two Lambda functions every hour will incur more charges than invoking one.

? Option C is incorrect because GuardDuty does not provide drift detection for CloudFormation stacks. GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior in AWS accounts and workloads. It does not monitor or report on configuration changes or drifts in CloudFormation.

stacks.

? Option D is correct because it leverages AWS Config and its managed rule for drift detection. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. It can detect configuration changes and drifts in CloudFormation stacks using the cloudformation-stack-drift-detection-check-managed-rule. This rule triggers an AWS Config event when a stack drifts from its expected template configuration. By creating a second EventBridge rule that reacts to this event for the specific stack, the DevOps engineer can configure the SNS topic as a target and receive a notification as soon as possible when drift is detected.

References:

? AWS Config

? Amazon SNS subscription filter policies

? Amazon EventBridge rules

NEW QUESTION 43

A company is launching an application. The application must use only approved AWS services. The account that runs the application was created less than 1 year ago and is assigned to an AWS Organizations OU.

The company needs to create a new Organizations account structure. The account structure must have an appropriate SCP that supports the use of only services that are currently active in the AWS account.

The company will use AWS Identity and Access Management (IAM) Access Analyzer in the solution.

Which solution will meet these requirements?

- A. Create an SCP that allows the services that IAM Access Analyzer identifies
- B. Create an OU for the account
- C. Move the account into the new OU
- D. Attach the new SCP to the new OU
- E. Detach the default FullAWSAccess SCP from the new OU.
- F. Create an SCP that denies the services that IAM Access Analyzer identifies
- G. Create an OU for the account
- H. Move the account into the new OU
- I. Attach the new SCP to the new OU.
- J. Create an SCP that allows the services that IAM Access Analyzer identifies
- K. Attach the new SCP to the organization's root.
- L. Create an SCP that allows the services that IAM Access Analyzer identifies
- M. Create an OU for the account
- N. Move the account into the new OU
- O. Attach the new SCP to the management account
- P. Detach the default FullAWSAccess SCP from the new OU.

Answer: A

Explanation:

To meet the requirements of creating a new Organizations account structure with an appropriate SCP that supports the use of only services that are currently active in the AWS account, the company should use the following solution:

? Create an SCP that allows the services that IAM Access Analyzer identifies. IAM Access Analyzer is a service that helps identify potential resource-access risks by analyzing resource-based policies in the AWS environment. IAM Access Analyzer can also generate IAM policies based on access activity in the AWS CloudTrail logs. By using IAM Access Analyzer, the company can create an SCP that grants only the permissions that are required for the application to run, and denies all other services. This way, the company can enforce the use of only approved AWS services and reduce the risk of unauthorized access¹²

? Create an OU for the account. Move the account into the new OU. An OU is a container for accounts within an organization that enables you to group accounts that have similar business or security requirements. By creating an OU for the account, the company can apply policies and manage settings for the account as a group. The company should move the account into the new OU to make it subject to the policies attached to the OU³

? Attach the new SCP to the new OU. Detach the default FullAWSAccess SCP from the new OU. An SCP is a type of policy that specifies the maximum permissions for an organization or organizational unit (OU). By attaching the new SCP to the new OU, the company can restrict the services that are available to all accounts in that OU, including the account that runs the application. The company should also detach the default FullAWSAccess SCP from the new OU, because this policy allows all actions on all AWS services and might override or conflict with the new SCP⁴⁵

The other options are not correct because they do not meet the requirements or follow best practices. Creating an SCP that denies the services that IAM Access Analyzer identifies is not a good option because it might not cover all possible services that are not approved or required for the application. A deny policy is also more difficult to maintain and update than an allow policy. Creating an SCP that allows the services that IAM Access Analyzer identifies and attaching it to the organization's root is not a good option because it might affect other accounts and OUs in the organization that have different service requirements or approvals. Creating an SCP that allows the services that IAM Access Analyzer identifies and attaching it to the management account is not a valid option because SCPs cannot be attached directly to accounts, only to OUs or roots.

References:

? 1: Using AWS Identity and Access Management Access Analyzer - AWS Identity and Access Management

? 2: Generate a policy based on access activity - AWS Identity and Access Management

? 3: Organizing your accounts into OUs - AWS Organizations

? 4: Service control policies - AWS Organizations

? 5: How SCPs work - AWS Organizations

NEW QUESTION 44

A company uses an Amazon API Gateway regional REST API to host its application API. The REST API has a custom domain. The REST API's default endpoint is deactivated.

The company's internal teams consume the API. The company wants to use mutual TLS between the API and the internal teams as an additional layer of authentication.

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

- A. Use AWS Certificate Manager (ACM) to create a private certificate authority (CA). Provision a client certificate that is signed by the private CA.
- B. Provision a client certificate that is signed by a public certificate authority (CA). Import the certificate into AWS Certificate Manager (ACM).
- C. Upload the provisioned client certificate to an Amazon S3 bucket
- D. Configure the API Gateway mutual TLS to use the client certificate that is stored in the S3 bucket as the trust store.
- E. Upload the provisioned client certificate private key to an Amazon S3 bucket
- F. Configure the API Gateway mutual TLS to use the private key that is stored in the S3 bucket as the trust store.
- G. Upload the root private certificate authority (CA) certificate to an Amazon S3 bucket
- H. Configure the API Gateway mutual TLS to use the private CA certificate that is stored in the S3 bucket as the trust store.

Answer: AE

Explanation:

Mutual TLS (mTLS) authentication requires two-way authentication between the client and the server. For Amazon API Gateway, you can enable mTLS for a custom domain name, which requires clients to present X.509 certificates to verify their identity to access your API. To set up mTLS, you would typically use AWS Certificate Manager (ACM) to create a private certificate authority (CA) and provision a client certificate signed by this private CA. The root CA certificate is then uploaded to an Amazon S3 bucket and configured in API Gateway as the trust store¹².

References:

? Introducing mutual TLS authentication for Amazon API Gateway¹.

? Configuring mutual TLS authentication for a REST API².

? AWS Private Certificate Authority details³.

? AWS Certificate Manager Private Certificate Authority updates⁴.

NEW QUESTION 45

A company has multiple accounts in an organization in AWS Organizations. The company's SecOps team needs to receive an Amazon Simple Notification Service (Amazon SNS) notification if any account in the organization turns off the Block Public Access feature on an Amazon S3 bucket. A DevOps engineer must implement this change without affecting the operation of any AWS accounts. The implementation must ensure that individual member accounts in the organization cannot turn off the notification.

Which solution will meet these requirements?

- A. Designate an account to be the delegated Amazon GuardDuty administrator account
- B. Turn on GuardDuty for all accounts across the organization
- C. In the GuardDuty administrator account, create an SNS topic
- D. Subscribe the SecOps team's email address to the SNS topic
- E. In the same account, create an Amazon EventBridge rule that uses an event pattern for GuardDuty findings and a target of the SNS topic.
- F. Create an AWS CloudFormation template that creates an SNS topic and subscribes the SecOps team's email address to the SNS topic
- G. In the template, include an Amazon EventBridge rule that uses an event pattern of CloudTrail activity for s3:PutBucketPublicAccessBlock and a target of the SNS topic
- H. Deploy the stack to every account in the organization by using CloudFormation StackSets.
- I. Turn on AWS Config across the organization
- J. In the delegated administrator account, create an SNS topic
- K. Subscribe the SecOps team's email address to the SNS topic
- L. Deploy a conformance pack that uses the s3-bucket-level-public-access-prohibited AWS Config managed rule in each account and uses an AWS Systems Manager document to publish an event to the SNS topic to notify the SecOps team.
- M. Turn on Amazon Inspector across the organization
- N. In the Amazon Inspector delegated administrator account, create an SNS topic
- O. Subscribe the SecOps team's email address to the SNS topic
- P. In the same account, create an Amazon EventBridge rule that uses an event pattern for public network exposure of the S3 bucket and publishes an event to the SNS topic to notify the SecOps team.

Answer: C

Explanation:

Amazon GuardDuty is primarily on threat detection and response, not configuration monitoring. A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a Region or across an organization in AWS Organizations.

<https://docs.aws.amazon.com/config/latest/developerguide/conformance-packs.html> <https://docs.aws.amazon.com/config/latest/developerguide/s3-account-level-public-access-blocks.html>

NEW QUESTION 46

A company manages AWS accounts for application teams in AWS Control Tower. Individual application teams are responsible for securing their respective AWS accounts.

A DevOps engineer needs to enable Amazon GuardDuty for all AWS accounts in which the application teams have not already enabled GuardDuty. The DevOps engineer is using AWS CloudFormation StackSets from the AWS Control Tower management account.

How should the DevOps engineer configure the CloudFormation template to prevent failure during the StackSets deployment?

- A. Create a CloudFormation custom resource that invokes an AWS Lambda function
- B. Configure the Lambda function to conditionally enable GuardDuty if GuardDuty is not already enabled in the accounts.
- C. Use the Conditions section of the CloudFormation template to enable GuardDuty in accounts where GuardDuty is not already enabled.
- D. Use the CloudFormation Fn::GetAtt intrinsic function to check whether GuardDuty is already enabled. If GuardDuty is not already enabled, use the Resources section of the CloudFormation template to enable GuardDuty.
- E. GetAtt intrinsic function to check whether GuardDuty is already enabled. If GuardDuty is not already enabled, use the Resources section of the CloudFormation template to enable GuardDuty.
- F. Manually discover the list of AWS account IDs where GuardDuty is not enabled. Use the CloudFormation Fn::ImportValue intrinsic function to import the list of account IDs into the CloudFormation template to skip deployment for the listed AWS accounts.

Answer: A

Explanation:

This solution will meet the requirements because it will use a CloudFormation custom resource to execute custom logic during the stack set operation. A custom resource is a resource that you define in your template and that is associated with an AWS Lambda function. The Lambda function runs whenever the custom resource is created, updated, or deleted, and can perform any actions that are supported by the AWS SDK. In this case, the Lambda function can use the GuardDuty API to check whether GuardDuty is already enabled in each target account, and if not, enable it. This way, the DevOps engineer can avoid deploying the stack set to accounts that already have GuardDuty enabled, and prevent failure during the deployment.

NEW QUESTION 48

A DevOps engineer is creating an AWS CloudFormation template to deploy a web service. The web service will run on Amazon EC2 instances in a private subnet behind an Application Load Balancer (ALB). The DevOps engineer must ensure that the service can accept requests from clients that have IPv6 addresses.

What should the DevOps engineer do with the CloudFormation template so that IPv6 clients can access the web service?

- A. Add an IPv6 CIDR block to the VPC and the private subnet for the EC2 instance
- B. Create route table entries for the IPv6 network, use EC2 instance types that support IPv6, and assign IPv6 addresses to each EC2 instance.

- C. Assign each EC2 instance an IPv6 Elastic IP address
- D. Create a target group, and add the EC2 instances as target
- E. Create a listener on port 443 of the ALB, and associate the target group with the ALB.
- F. Replace the ALB with a Network Load Balancer (NLB). Add an IPv6 CIDR block to the VPC and subnets for the NLB, and assign the NLB an IPv6 Elastic IP address.
- G. Add an IPv6 CIDR block to the VPC and subnets for the AL
- H. Create a listener on port 443. and specify the dualstack IP address type on the AL
- I. Create a target group, and add the EC2 instances as target
- J. Associate the target group with the ALB.

Answer: D

Explanation:

it involves adding an IPv6 CIDR block to the VPC and subnets for the ALB and specifying the dualstack IP address type on the ALB listener. This allows the ALB to listen on both IPv4 and IPv6 addresses, and forward requests to the EC2 instances that are added as targets to the target group associated with the ALB.

NEW QUESTION 49

A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps engineer is tasked with minimizing application response times and improving availability for users in both Regions.

Which combination of actions should be taken to address the latency issues? (Choose three.)

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

Answer: CDE

Explanation:

C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group. This will allow users in the new Region to access the application with lower latency by reducing the network hops between the user and the application servers.

* D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB. This will enable Route 53 to route user traffic to the nearest healthy ALB, based on the latency between the user and the ALBs.

* F. Convert the DynamoDB table to a global table. This will enable reads and writes to the table in both Regions with low latency, improving the overall response time of the application

NEW QUESTION 53

A DevOps engineer needs to back up sensitive Amazon S3 objects that are stored within an S3 bucket with a private bucket policy using S3 cross-Region replication functionality. The objects need to be copied to a target bucket in a different AWS Region and account.

Which combination of actions should be performed to enable this replication? (Choose three.)

- A. Create a replication IAM role in the source account
- B. Create a replication IAM role in the target account.
- C. Add statements to the source bucket policy allowing the replication IAM role to replicate objects.
- D. Add statements to the target bucket policy allowing the replication IAM role to replicate objects.
- E. Create a replication rule in the source bucket to enable the replication.
- F. Create a replication rule in the target bucket to enable the replication.

Answer: ADE

Explanation:

S3 cross-Region replication (CRR) automatically replicates data between buckets across different AWS Regions. To enable CRR, you need to add a replication configuration to your source bucket that specifies the destination bucket, the IAM role, and the encryption type (optional). You also need to grant permissions to the IAM role to perform replication actions on both the source and destination buckets. Additionally, you can choose the destination storage class and enable additional replication options such as S3 Replication Time Control (S3 RTC) or S3 Batch Replication. <https://medium.com/cloud-techies/s3-same-region-replication-srr-and-cross-region-replication-crr-34d446806bab> <https://aws.amazon.com/getting-started/hands-on/replicate-data-using-amazon-s3-replication/> <https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html>

NEW QUESTION 54

An Amazon EC2 instance is running in a VPC and needs to download an object from a restricted Amazon S3 bucket. When the DevOps engineer tries to download the object, an AccessDenied error is received,

What are the possible causes for this error? (Select TWO,)

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

Answer: BD

Explanation:

These are the possible causes for the AccessDenied error because they affect the permissions to access the S3 object from the EC2 instance. An S3 bucket policy is a resource-based policy that defines who can access the bucket and its objects, and what actions they can perform. An IAM role is an identity that can be assumed by an EC2 instance to grant it permissions to access AWS services and resources. If there is an error in the S3 bucket policy or the IAM role configuration, such as a missing or incorrect statement, condition, or principal, then the EC2 instance may not have the necessary permissions to download the

object from the S3 bucket . <https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies.html>
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 59

A company has multiple AWS accounts. The company uses AWS IAM Identity Center (AWS Single Sign-On) that is integrated with AWS Toolkit for Microsoft Azure DevOps. The attributes for access control feature is enabled in IAM Identity Center.

The attribute mapping list contains two entries. The department key is mapped to

`${path:enterprise.department}`. The costCenter key is mapped to

`${path:enterprise.costCenter}`.

All existing Amazon EC2 instances have a department tag that corresponds to three company departments (d1, d2, d3). A DevOps engineer must create policies based on the matching attributes. The policies must minimize administrative effort and must grant each Azure AD user access to only the EC2 instances that are tagged with the user's respective department name.

Which condition key should the DevOps engineer include in the custom permissions policies to meet these requirements?

A.

```
"Condition": {
  "ForAllValues:StringEquals": {
    "aws:TagKeys": ["department"]
  }
}
```

B.

```
"Condition": {
  "StringEquals": {
    "aws:PrincipalTag/department": "${aws:ResourceTag/department}"
  }
}
```

C.

```
"Condition": {
  "StringEquals": {
    "ec2:ResourceTag/department": "${aws:PrincipalTag/department}"
  }
}
```

D.

```
"Condition": {
  "ForAllValues:StringEquals": {
    "ec2:ResourceTag/department": ["d1", "d2", "d3"]
  }
}
```

A.

Answer: C

Explanation:

<https://docs.aws.amazon.com/singlesignon/latest/userguide/configure-abac.html>

NEW QUESTION 62

A DevOps team is merging code revisions for an application that uses an Amazon RDS Multi-AZ DB cluster for its production database. The DevOps team uses continuous integration to periodically verify that the application works. The DevOps team needs to test the changes before the changes are deployed to the production database.

Which solution will meet these requirements'?

A. Use a buildspec file in AWS CodeBuild to restore the DB cluster from a snapshot of the production database run integration tests, and drop the restored database after verification.

B. Deploy the application to productio

C. Configure an audit log of data control language (DCL) operations to capture database activities to perform if verification fails.

D. Create a snapshot of the DB duster before deploying the application Use the Update requires Replacement property on the DB instance in AWS CloudFormation to deploy the application and apply the changes.

E. Ensure that the DB cluster is a Multi-AZ deploymen

F. Deploy the application with the update

G. Fail over to the standby instance if verification fails.

Answer: A

Explanation:

This solution will meet the requirements because it will create a temporary copy of the production database using a snapshot, run the integration tests on the copy, and delete the copy after the tests are done. This way, the production database will not be affected by the code revisions, and the DevOps team can test the changes before deploying them to production. A buildspec file is a YAML file that contains the commands and settings that CodeBuild uses to run a build¹. The buildspec file can specify the steps to restore the DB cluster from a snapshot, run the integration tests, and drop the restored database².

NEW QUESTION 64

A company's DevOps engineer is creating an AWS Lambda function to process notifications from an Amazon Simple Notification Service (Amazon SNS) topic. The Lambda function will process the notification messages and will write the contents of the notification messages to an Amazon RDS Multi-AZ DB instance. During testing a database administrator accidentally shut down the DB instance. While the database was down the company lost several of the SNS notification messages that were delivered during that time.

The DevOps engineer needs to prevent the loss of notification messages in the future Which solutions will meet this requirement? (Select TWO.)

- A. Replace the RDS Multi-AZ DB instance with an Amazon DynamoDB table.
- B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination of the Lambda function.
- C. Configure an Amazon Simple Queue Service (Amazon SQS) dead-letter queue for the SNS topic.
- D. Subscribe an Amazon Simple Queue Service (Amazon SQS) queue to the SNS topic Configure the Lambda function to process messages from the SQS queue.
- E. Replace the SNS topic with an Amazon EventBridge event bus Configure an EventBridge rule on the new event bus to invoke the Lambda function for each event.

Answer: CD

Explanation:

These solutions will meet the requirement because they will prevent the loss of notification messages in the future. An Amazon SQS queue is a service that provides a reliable, scalable, and secure message queue for asynchronous communication between distributed components. You can use an SQS queue to buffer messages from an SNS topic and ensure that they are delivered and processed by a Lambda function, even if the function or the database is temporarily unavailable.

Option C will configure an SQS dead-letter queue for the SNS topic. A dead-letter queue is a queue that receives messages that could not be delivered to any subscriber after a specified number of retries. You can use a dead-letter queue to store and analyze failed messages, or to reprocess them later. This way, you can avoid losing messages that could not be delivered to the Lambda function due to network errors, throttling, or other issues. Option D will subscribe an SQS queue to the SNS topic and configure the Lambda function to process messages from the SQS queue. This will decouple the SNS topic from the Lambda function and provide more flexibility and control over the message delivery and processing. You can use an SQS queue to store messages from the SNS topic until they are ready to be processed by the Lambda function, and also to retry processing in case of failures. This way, you can avoid losing messages that could not be processed by the Lambda function due to database errors, timeouts, or other issues.

NEW QUESTION 66

A company requires its developers to tag all Amazon Elastic Block Store (Amazon EBS) volumes in an account to indicate a desired backup frequency. This requirement Includes EBS volumes that do not require backups. The company uses custom tags named Backup_Frequency that have values of none, daily, or weekly that correspond to the desired backup frequency. An audit finds that developers are occasionally not tagging the EBS volumes.

A DevOps engineer needs to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified.

Which solution will meet these requirements?

- A. Set up AWS Config in the account
- B. Create a custom rule that returns a compliance failure for all Amazon EC2 resources that do not have a Backup Frequency tag applied. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
- C. Set up AWS Config in the account
- D. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied
- E. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
- F. Turn on AWS CloudTrail in the account
- G. Create an Amazon EventBridge rule that reacts to EBS CreateVolume event
- H. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly
- I. Specify the runbook as the target of the rule.
- J. Turn on AWS CloudTrail in the account
- K. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events or EBS ModifyVolume event
- L. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly
- M. Specify the runbook as the target of the rule.

Answer: B

Explanation:

The following are the steps that the DevOps engineer should take to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified:

? Set up AWS Config in the account.

? Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied.

? Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.

The managed rule AWS::Config::EBSVolumesWithoutBackupTag will return a compliance failure for any EBS volume that does not have the Backup_Frequency tag applied. The remediation action will then use the Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly to the EBS volume.

NEW QUESTION 67

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

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/ASGettingNotifications.html#auto-scaling-sns-notifications>

NEW QUESTION 69

A company's security policies require the use of security hardened AMIS in production environments. A DevOps engineer has used EC2 Image Builder to create a pipeline that builds the AMIs on a recurring schedule.

The DevOps engineer needs to update the launch templates of the company's Auto Scaling groups. The Auto Scaling groups must use the newest AMIS during the launch of Amazon EC2 instances.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Configure an Amazon EventBridge rule to receive new AMI events from Image Builder
- B. Target an AWS Systems Manager Run Command document that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
- C. Configure an Amazon EventBridge rule to receive new AMI events from Image Builder
- D. Target an AWS Lambda function that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
- E. Configure the launch template to use a value from AWS Systems Manager Parameter Store for the AMI ID
- F. Configure the Image Builder pipeline to update the Parameter Store value with the newest AMI ID.
- G. Configure the Image Builder distribution settings to update the launch templates with the newest AMI ID
- H. Configure the Auto Scaling groups to use the newest version of the launch template.

Answer: C

Explanation:

? The most operationally efficient solution is to use AWS Systems Manager Parameter Store¹ to store the AMI ID and reference it in the launch template². This way, the launch template does not need to be updated every time a new AMI is created by Image Builder. Instead, the Image Builder pipeline can update the Parameter Store value with the newest AMI ID³, and the Auto Scaling group can launch instances using the latest value from Parameter Store.

? The other solutions require updating the launch template or creating a new version of it every time a new AMI is created, which adds complexity and overhead. Additionally, using EventBridge rules and Lambda functions or Run Command documents introduces additional dependencies and potential points of failure.

References: 1: AWS Systems Manager Parameter Store 2: Using AWS Systems Manager parameters instead of AMI IDs in launch templates 3: Update an SSM parameter with Image Builder

NEW QUESTION 71

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 database population script.
- C. Modify the S3 bucket settings to enable HTTPS basic authentication and specify a token
- 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 policy
- F. Modify the service role for the CodeBuild project to include Amazon S3 access
- G. Use the AWS CLI to download the database population script.
- H. Remove unauthenticated access from the S3 bucket with a bucket policy
- I. Use the AWS CLI to download the database population script using an IAM access key and a secret access key.

Answer: C

Explanation:

A bucket policy is a resource-based policy that defines who can access a specific S3 bucket and what actions they can perform on it. By removing unauthenticated access from the bucket policy, you can prevent anyone without valid credentials from accessing the bucket. A service role is an IAM role that allows an AWS service, such as CodeBuild, to perform actions on your behalf. By modifying the service role for the CodeBuild project to include Amazon S3 access, you can grant the project permission to read and write objects in the S3 bucket. The AWS CLI is a command-line tool that allows you to interact with AWS services, such as S3, using commands in your terminal. By using the AWS CLI to download the database population script, you can leverage the service role credentials and encryption to secure the data transfer.

For more information, you can refer to these web pages:

? [Using bucket policies and user policies - Amazon Simple Storage Service]

? [Create a service role for CodeBuild - AWS CodeBuild]

? [AWS Command Line Interface]

NEW QUESTION 75

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


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

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

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

Answer: BCE

Explanation:

B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable. C. Store the DB_PASSWORD as a SecureString value in AWS Systems Manager Parameter Store and then remove the DB_PASSWORD from the environment variables. E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.

NEW QUESTION 79

A company uses AWS Organizations to manage its AWS accounts. The company has a root OU that has a child OU. The root OU has an SCP that allows all actions on all resources. The child OU has an SCP that allows all actions for Amazon DynamoDB and AWS Lambda, and denies all other actions. The company has an AWS account that is named vendor-data in the child OU. A DevOps engineer has an 1AM user that is attached to the AdministratorAccess 1AM policy in the vendor-data account. The DevOps engineer attempts to launch an Amazon EC2 instance in the vendor-data account but receives an access denied error.

Which change should the DevOps engineer make to launch the EC2 instance in the vendor-data account?

- A. Attach the AmazonEC2FullAccess 1AM policy to the 1AM user.
- B. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the vendor-data account.
- C. Update the SCP in the child OU to allow all actions for Amazon EC2.
- D. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the root OU.

Answer: C

Explanation:

The correct answer is C. Updating the SCP in the child OU to allow all actions for Amazon EC2 will enable the DevOps engineer to launch the EC2 instance in the vendor-data account. SCPs are applied to OUs and accounts in a hierarchical manner, meaning that the SCPs attached to the parent OU are inherited by the child OU and accounts. Therefore, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. By adding EC2 to the allowed actions in the child OU's SCP, the DevOps engineer can access EC2 resources in the vendor-data account.

Option A is incorrect because attaching the AmazonEC2FullAccess IAM policy to the IAM user will not grant the user access to EC2 resources. IAM policies are evaluated after SCPs, so even if the IAM policy allows EC2 actions, the SCP will still deny them.

Option B is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the vendor-data account will not work. SCPs are not cumulative, meaning that only one SCP is applied to an account at a time. The SCP attached to the account will be the SCP attached to the OU that contains the account. Therefore, option B will not change the SCP that is applied to the vendor-data account.

Option D is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the root OU will not work. As explained earlier, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. Therefore, option D will not affect the SCP that is applied to the vendor-data account.

NEW QUESTION 82

A company runs its container workloads in AWS App Runner. A DevOps engineer manages the company's container repository in Amazon Elastic Container Registry (Amazon ECR).

The DevOps engineer must implement a solution that continuously monitors the container repository. The solution must create a new container image when the solution detects an operating system vulnerability or language package vulnerability.

Which solution will meet these requirements?

- A. Use EC2 Image Builder to create a container image pipeline
- B. Use Amazon ECR as the target repository
- C. Turn on enhanced scanning on the ECR repository
- D. Create an Amazon EventBridge rule to capture an Inspector2 finding even
- E. Use the event to invoke the image pipeline
- F. Re-upload the container to the repository.
- G. Use EC2 Image Builder to create a container image pipeline
- H. Use Amazon ECR as the target repository
- I. Enable Amazon GuardDuty Malware Protection on the container workload
- J. Create an Amazon EventBridge rule to capture a GuardDuty finding even
- K. Use the event to invoke the image pipeline.
- L. Create an AWS CodeBuild project to create a container image
- M. Use Amazon ECR as the target repository

- N. Turn on basic scanning on the repositior
- O. Create an Amazon EventBridge rule to capture an ECR image action even
- P. Use the event to invoke the CodeBuild projec
- Q. Re-upload the container to the repository.
- R. Create an AWS CodeBuild project to create a container imag
- S. Use Amazon ECR as the target repositior
- T. Configure AWS Systems Manager Compliance to scan all managed node
- . Create an Amazon EventBridge rule to capture a configuration compliance state change even
- . Use the event to invoke the CodeBuild project.

Answer: A

Explanation:

The solution that meets the requirements is to use EC2 Image Builder to create a container image pipeline, use Amazon ECR as the target repository, turn on enhanced scanning on the ECR repository, create an Amazon EventBridge rule to capture an Inspector2 finding event, and use the event to invoke the image pipeline. Re-upload the container to the repository.

This solution will continuously monitor the container repository for vulnerabilities using enhanced scanning, which is a feature of Amazon ECR that provides detailed information and guidance on how to fix security issues found in your container images. Enhanced scanning uses Inspector2, a security assessment service that integrates with Amazon ECR and generates findings for any vulnerabilities detected in your images. You can use Amazon EventBridge to create a rule that triggers an action when an Inspector2 finding event occurs. The action can be to invoke an EC2 Image Builder pipeline, which is a service that automates the creation of container images. The pipeline can use the latest patches and updates to build a new container image and upload it to the same ECR repository, replacing the vulnerable image.

The other options are not correct because they do not meet all the requirements or use services that are not relevant for the scenario.

Option B is not correct because it uses Amazon GuardDuty Malware Protection, which is a feature of GuardDuty that detects malicious activity and unauthorized behavior on your AWS accounts and resources. GuardDuty does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder.

Option C is not correct because it uses basic scanning on the ECR repository, which only provides a summary of the vulnerabilities found in your container images. Basic scanning does not use Inspector2 or generate findings that can be captured by Amazon EventBridge. Moreover, basic scanning does not provide guidance on how to fix the vulnerabilities.

Option D is not correct because it uses AWS Systems Manager Compliance, which is a feature of Systems Manager that helps you monitor and manage the compliance status of your AWS resources based on AWS Config rules and AWS Security Hub standards. Systems Manager Compliance does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder.

NEW QUESTION 85

A company hosts applications in its AWS account Each application logs to an individual Amazon CloudWatch log group. The company's CloudWatch costs for ingestion are increasing

A DevOps engineer needs to Identify which applications are the source of the increased logging costs.

Which solution Will meet these requirements?

- A. Use CloudWatch metrics to create a custom expression that Identifies the CloudWatch log groups that have the most data being written to them.
- B. Use CloudWatch Logs Insights to create a set of queries for the application log groups to Identify the number of logs written for a period of time
- C. Use AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage
- D. Use AWS CloudTrail to filter for CreateLogStream events for each application

Answer: C

Explanation:

The correct answer is C.

A comprehensive and detailed explanation is:

? Option A is incorrect because using CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them is not a valid solution. CloudWatch metrics do not provide information about the size or volume of data being ingested by CloudWatch logs.

CloudWatch metrics only provide information about the number of events, bytes, and errors that occur within a log group or stream. Moreover, creating a custom expression with CloudWatch metrics would require using the search_web tool, which is not necessary for this use case.

? Option B is incorrect because using CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time is not a valid solution. CloudWatch Logs Insights can help analyze and filter log events based on patterns and expressions, but it does not provide information about the cost or billing of CloudWatch logs. CloudWatch Logs Insights also charges based on the amount of data scanned by each query, which could increase the logging costs further.

? Option C is correct because using AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage is a valid solution. AWS Cost Explorer is a tool that helps visualize, understand, and manage AWS costs and usage over time. AWS Cost Explorer can generate custom reports that show the breakdown of costs by service, region, account, tag, or any other dimension. AWS Cost Explorer can also filter and group costs by usage type, which can help identify the specific CloudWatch log groups that are the source of the increased logging costs.

? Option D is incorrect because using AWS CloudTrail to filter for CreateLogStream events for each application is not a valid solution. AWS CloudTrail is a service that records API calls and account activity for AWS services, including CloudWatch logs. However, AWS CloudTrail does not provide information about the cost or billing of CloudWatch logs. Filtering for CreateLogStream events would only show when a new log stream was created within a log group, but not how much data was ingested or stored by that log stream.

References:

? CloudWatch Metrics

? CloudWatch Logs Insights

? AWS Cost Explorer

? AWS CloudTrail

NEW QUESTION 88

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 endpoint
- B. Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to invoke an AWS Lambda function that will promote the replica instance as the primary.
- C. Create an Aurora custom endpoint to point to the primary database instanc

- D. Configure the application to use this endpoint
- E. 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.
- F. 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 point the application to the newly promoted instance
- G. Create an Amazon CloudWatch alarm to invoke this Lambda function after the failure event occurs.
- H. Store the Aurora endpoint in AWS Systems Manager Parameter Store
- I. Create an Amazon EventBridge 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
- J. Code the application to reload the endpoint from Parameter Store if a database connection fails.

Answer: D

Explanation:

EventBridge is needed to detect the database failure. Lambda is needed to promote the replica as it's in another Region (manual promotion, otherwise). Storing and updating the endpoint in Parameter store is important in updating the application. Look at High Availability section of Aurora FAQ:
<https://aws.amazon.com/rds/aurora/faqs/>

NEW QUESTION 93

To run an application, a DevOps engineer launches an Amazon EC2 instance 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 attached
- B. Once the application is installed and running, run a script to disassociate the Elastic IP addresses afterwards.
- C. Set up a NAT gateway
- D. Deploy the EC2 instances to a private subnet
- 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 allow only outbound traffic to the artifact repository
- 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 96

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 Logs. Use CloudWatch Logs Insights to query both sets of logs.
- C. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon Kinesis. Configure AWS CloudTrail to deliver the API logs to Kinesis. Use Kinesis to load the data into Amazon Redshift. Use Amazon Redshift to query both sets of logs.
- D. 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: D

Explanation:

This solution will meet the requirements because it will use Amazon S3 as a common data lake for both the application logs and the API logs. Amazon S3 is a service that provides scalable, durable, and secure object storage for any type of data. You can use the Amazon CloudWatch agent to send logs from your EC2 instances to S3 buckets, and use AWS CloudTrail to deliver the API logs to S3 buckets as well. You can also use Amazon Athena to query both sets of logs in S3 using standard SQL, without loading or transforming them. Athena is a serverless interactive query service that allows you to analyze data in S3 using a variety of data formats, such as JSON, CSV, Parquet, and ORC.

NEW QUESTION 100

A rapidly growing company wants to scale for developer demand for AWS development environments. Development environments are created manually in the AWS Management Console. The networking team uses AWS CloudFormation to manage the networking infrastructure, exporting stack output values for the Amazon VPC and all subnets. The development environments have common standards, such as Application Load Balancers, Amazon EC2 Auto Scaling groups, security groups, and Amazon DynamoDB tables.

To keep up with demand, the DevOps engineer wants to automate the creation of development environments. Because the infrastructure required to support the application is expected to grow, there must be a way to easily update the deployed infrastructure. CloudFormation will be used to create a template for the development environments.

Which approach will meet these requirements and quickly provide consistent AWS environments for developers?

- A. Use Fn::ImportValue intrinsic functions in the Resources section of the template to retrieve Virtual Private Cloud (VPC) and subnet values
- B. Use CloudFormation StackSets for the development environments, using the Count input parameter to indicate the number of environments needed
- C. Use the UpdateStackSet command to update existing development environments.
- D. Use nested stacks to define common infrastructure components
- E. To access the exported values, use TemplateURL to reference the networking team's template
- F. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root template
- G. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.

- H. Use nested stacks to define common infrastructure component
- I. Use Fn::ImportValue intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet value
- J. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
- K. Use Fn::ImportValue intrinsic functions in the Parameters section of the root template to retrieve Virtual Private Cloud (VPC) and subnet value
- L. Define the development resources in the order they need to be created in the CloudFormation nested stack
- M. Use the CreateChangeSet
- N. and ExecuteChangeSet commands to update existing development environments.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-importvalue.html>
<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-importvalue.html> CF of network exports the VPC, subnet or needed information CF of application imports the above information to its stack and UpdateChangeSet/ ExecuteChangeSet

NEW QUESTION 102

A company uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups.

The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account.

When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account's backup vault.

Which combination of steps must the company take so that backups can be copied to the new account's backup vault? (Select TWO.)

- A. Edit the backup vault access policy in the new account to allow access to the primary account.
- B. Edit the backup vault access policy in the primary account to allow access to the new account.
- C. Edit the backup vault access policy in the primary account to allow access to the KMS key in the new account.
- D. Edit the key policy of the KMS key in the primary account to share the key with the new account.
- E. Edit the key policy of the KMS key in the new account to share the key with the primary account.

Answer: AE

Explanation:

To enable cross-account backup, the company needs to grant permissions to both the backup vault and the KMS key in the destination account. The backup vault access policy in the destination account must allow the primary account to copy backups into the vault. The key policy of the KMS key in the destination account must allow the primary account to use the key to encrypt and decrypt the backups. These steps are described in the AWS documentation¹². Therefore, the correct answer is A and E.

References:

? 1: Creating backup copies across AWS accounts - AWS Backup

? 2: Using AWS Backup with AWS Organizations - AWS Backup

NEW QUESTION 105

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

How should a DevOps engineer meet these requirements?

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

Answer: D

NEW QUESTION 106

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

Answer: B

Explanation:

<https://aws.amazon.com/blogs/security/how-to-monitor-and-visualize-failed-ssh-access-attempts-to-amazon-ec2-linux-instances/>

NEW QUESTION 110

A company manages multiple AWS accounts in AWS Organizations. The company's security policy states that AWS account root user credentials for member accounts must not be used. The company monitors access to the root user credentials.

A recent alert shows that the root user in a member account launched an Amazon EC2 instance. A DevOps engineer must create an SCP at the organization's root level that will prevent the root user in member accounts from making any AWS service API calls.

Which SCP will meet these requirements?

A)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringNotLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
      }
    }
  ]
}
```

B)

```
"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Deny",
    "Action": "*",
    "Resource": "*",
    "Principal": { "AWS": "arn:aws:iam::*:root" }
  }
]
```

C)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
      }
    }
  ]
}
```

D)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Principal": "root"
    }
  ]
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

NEW QUESTION 114

An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to external identity provider (IdP) and has configured SAML 2.0. The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources.

Which combination of steps will meet these requirements? (Choose three.)

- A. Create IAM policies that include the required permission
- B. Include the aws:PrincipalTag condition key.
- C. Create permission set
- D. Attach an inline policy that includes the required permissions and uses the aws:PrincipalTag condition key to scope the permissions.
- E. Create a group in the Id
- F. Place users in the grou
- G. Assign the group to accounts and the permission sets in IAM Identity Center.
- H. Create a group in the Id
- I. Place users in the grou
- J. Assign the group to OUs and IAM policies.
- K. Enable attributes for access control in IAM Identity Cente
- L. Apply tags to user
- M. Map the tags as key-value pairs.
- N. Enable attributes for access control in IAM Identity Cente
- O. Map attributes from the IdP as key-value pairs.

Answer: BCF

Explanation:

Using the principalTag in the Permission Set inline policy a logged in user belonging to a specific AD group in the IDP can be permitted access to perform operations on certain resources if their group matches the group used in the PrincipleTag. Basically you are narrowing the scope of privileges assigned via Permission policies conditionally based on whether the logged in user belongs to a specific AD Group in IDP. The mapping of the AD group to the request attributes can be done using SSO attributes where we can pass other attributes like the SAML token as well.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/abac.html>

NEW QUESTION 119

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

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

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

Answer: A

Explanation:

You can run custom recipes manually, but the best approach is usually to have AWS OpsWorks Stacks run them automatically. Every layer has a set of built-in recipes assigned each of five lifecycle events—Setup, Configure, Deploy, Undeploy, and Shutdown. Each time an event occurs for an instance, AWS OpsWorks Stacks runs the associated recipes for each of the instance's layers, which handle the corresponding tasks. For example, when an instance finishes booting, AWS OpsWorks Stacks triggers a Setup event. This event runs the associated layer's Setup recipes, which typically handle tasks such as installing and configuring packages

NEW QUESTION 121

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 126

A company uses AWS Secrets Manager to store a set of sensitive API keys that an AWS Lambda function uses. When the Lambda function is invoked, the Lambda function retrieves the API keys and makes an API call to an external service. The Secrets Manager secret is encrypted with the default AWS Key Management Service (AWS KMS) key.

A DevOps engineer needs to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege.

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

- A. Update the default KMS key for Secrets Manager to allow only the Lambda function's execution role to decrypt.
- B. Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt
- C. Update Secrets Manager to use the new customer managed key.
- D. Create a KMS customer managed key that trusts Secrets Manager and allows the account's :root principal to decrypt
- E. Update Secrets Manager to use the new customer managed key.
- F. Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource level
- G. Configure the permissions so that the KMS key can encrypt the Secrets Manager secret.
- H. Remove all KMS permissions from the Lambda function's execution role.

Answer: BD

Explanation:

The requirement is to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege, which means granting the minimum permissions necessary to perform a task.

To do this, the DevOps engineer needs to use the following steps:

? Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt. A customer managed key is a symmetric encryption key that is fully managed by the customer. The customer can define the key policy, which specifies who can use and manage the key. By creating a customer managed key, the DevOps engineer can restrict the decryption permission to only the Lambda function's execution role, and prevent other principals from accessing the secret values. The customer managed key also needs to trust Secrets Manager, which means allowing Secrets Manager to use the key to encrypt and decrypt secrets on behalf of the customer.

? Update Secrets Manager to use the new customer managed key. Secrets Manager allows customers to choose which KMS key to use for encrypting each secret. By default, Secrets Manager uses the default KMS key for Secrets Manager, which is a service-managed key that is shared by all customers in the same AWS Region. By updating Secrets Manager to use the new customer managed key, the DevOps engineer can ensure that only the Lambda function's execution role can decrypt the secret values using that key.

? Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource level. The Lambda function's execution role is an IAM role that grants permissions to the Lambda function to access AWS services and resources. The role needs to have KMS permissions to use the customer managed key for decryption. However, to apply the principle of least privilege, the role should have the permissions scoped on the resource level, which means specifying the ARN of the customer managed key as a condition in the IAM policy statement. This way, the role can only use that specific key and not any other KMS keys in the account.

NEW QUESTION 131

A company uses AWS Organizations to manage multiple accounts. Information security policies require that all unencrypted Amazon EBS volumes be marked as non-compliant. A DevOps engineer needs to automatically deploy the solution and ensure that this compliance check is always present.

Which solution will accomplish this?

- A. Create an AWS CloudFormation template that defines an AWS Inspector rule to check whether EBS encryption is enable
- B. Save the template to an Amazon S3 bucket that has been shared with all accounts within the compan
- C. Update the account creation script pointing to the CloudFormation template in Amazon S3.
- D. Create an AWS Config organizational rule to check whether EBS encryption is enabled and deploy the rule using the AWS CL
- E. Create and apply an SCP to prohibit stopping and deleting AWS Config across the organization.
- F. Create an SCP in Organization
- G. Set the policy to prevent the launch of Amazon EC2 instances without encryption on the EBS volumes using a conditional expressio
- H. Apply the SCP to all AWS account
- I. Use Amazon Athena to analyze the AWS CloudTrail output, looking for events that deny an ec2: RunInstances action.
- J. Deploy an IAM role to all accounts from a single trusted account
- K. Build a pipeline with AWS CodePipeline with a stage in AWS Lambda to assume the IAM role, and list all EBS volumes in the account
- L. Publish a report to Amazon S3.

Answer: B

Explanation:

<https://docs.aws.amazon.com/config/latest/developerguide/ec2-ebs-encryption-by-default.html>

NEW QUESTION 136

.....

Relate Links

100% Pass Your DOP-C02 Exam with ExamBible Prep Materials

<https://www.exambible.com/DOP-C02-exam/>

Contact us

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

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