

Amazon.DOP-C02-KR.v2026-05-20.q183

□□□□:	DOP-C02-KR
□□□□:	AWS Certified DevOps Engineer - Professional (DOP-C02 Korean Version)
□□□:	Amazon
□□ □□ □□□:	183
□□:	v2026-05-20
# □□ □:	166
# □□ □□□:	1830
https://www.krdump.com/Amazon.DOP-C02-KR.v2026-05-20.q183.html	

NEW QUESTION: 1

□ □□□ AWS Organizations, AWS Control Tower, AWS Config, Terraform □ □□□□ AWS □□□ □□□□ □□□□□. □□□ □□□ □ □□ AWS □□□ VPC □ □□□ AWS Lambda □□□ □□□□□ □□ □□□. □□ □□□□ □□□□ □□ □□□□ □□□ □□ □□□ □□□ □ □□□□?

- A. AWS Control Tower □ □□□□ □□ □□□ □□(□□□□) □ □□□□□. □□ OU □□□ Lambda □ AWS CloudFormation □□□ □□ □ □□□ □□□ □□□□□□.
- B. □□□□ □□ □□ □□:Vpclds □□ □□ □□□□ □ SCP □ □□□□.
- C. VPC □ □□□□ □□ Lambda □□□ □□□□ □□ □□□ □□ AWS Config □□□ □□□□.
- D. lambda:Vpclds □ null □ □□ Lambda □□□ □□□□ □□□□□ □ SCP □ □□□□.

Answer: (SHOW ANSWER)

Use a Service Control Policy (SCP) with a Null condition on lambda:Vpclds to deny Lambda function creation or update when not VPC-attached. This enforces compliance across all accounts automatically without manual remediation, aligning with AWS Control Tower governance recommendations.

NEW QUESTION: 2

DevOps □□□□□ AWS Lambda □□□ □□□□□□□. Lambda □□□ □□ CloudFormation □□□ □□ □□□□ □□ □□□□□ AWS CloudFormation □□□□ □□ □□□ □□□□□□. □□ □□ Lambda □□□ □□ □□□ □□□□□. DevOps □□□□□ Lambda □□□ □□□□ Amazon EventBridge □□ □□□ □□□□□□□. □□□. Amazon Simple □□ □□□(Amazon SNS) □□□ AWS □□□ □□ □□□□□. DevOps □□□□□ □□□ □□ □□ SNS □□□ □□□□□□□. DevOps □□□□□ □ □□ □□ □□□□ □□□□□ □□□□ □□□ □ □□ □□□ □□□ □□□.

- A. SNS □□□ □□□□ □□□□□ □□ EventBridge □□□ □□□□□. Cloud Formation □□□ □□□□□ SNS □□ □□ □□□ □ □□□□. SNS tomc □ □□ □□ □□□ □□□□□.
- B. □ □□ Lambda □□□ □□□□ □□□ □□□□ □□ □□□ CloudFormation API □ □□□□□. SNS □□□ □□□□ □□□□□ □ □□ Lambda □□□ □□□□□. □□□□□ □□□□ □ □□ Lambda □ □□□□ □□□□□ □□ EventBridge □□□ □□□□□. □□

C. CloudFormation stacks are monitored by Amazon GuardDuty. CloudFormation stacks are monitored by GuardDuty. EventBridge rules can filter events based on the message body or attributes of the target service. SNS topics can be used as targets for EventBridge rules.

D. AWS Config provides drift detection for CloudFormation stacks. CloudFormation stacks are monitored by AWS Config. EventBridge rules can filter events based on the message body or attributes of the target service. SNS topics can be used as targets for EventBridge rules.

Answer: D (LEAVE A REPLY)

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: 3

A DevOps engineer is using AWS Config to monitor EC2 instances for attached instance profiles. The engineer wants to receive a notification when an instance profile is attached to an EC2 instance. Which of the following configurations will accomplish this goal?

A. RunInstances API, EventBridge rule, Lambda function, SNS topic.

B. ec2-instance-profile-attached managed rule, AWS Config, Systems Manager Automation runbook, SNS topic.

C. StartInstances API, EventBridge rule, Systems Manager Automation runbook.

D. AWS Config iam-role-managed-policy-check managed rule, Lambda function, SNS topic.

Answer: (SHOW ANSWER)

* AWS Config's ec2-instance-profile-attached managed rule checks for attached instance profiles.

* Config supports automatic remediation via Systems Manager Automation runbooks.

* This provides continuous compliance with minimal operational overhead.

* EventBridge and Lambda (A) require custom coding and risk missing existing instances.

* StartInstances (C) does not cover RunInstances and new instances.

* IAM-role managed policy check (D) does not check instance profile attachments.

References:

AWS Config Managed Rules

Config Automatic Remediation

NEW QUESTION: 4

□□ □□ □□□ □□ □□□□ □□ □□ □□□□ □□□□. □□ □□□ □□□ Amazon API Gateway□ □□ □□□□ Amazon Simple Queue Service(Amazon SQS) □□ □□□□ □□ □□□□□□. □□□ □□ □□□□□□□□ □□□□ □□□□ □□□□ □□ □□□□ □□□□.

□□□ □□□□ □□□□ □□□□ □□ □□□□□□□□ □□□ □□□□ □□□ □ □□□□. □□□ □□ □□□□ SQS □□□□ □□ □□□□. DevOps □□□□□□ □□□ □□□□ □□□□ □□ □ □□ □□□ □□ □□□□ □□□ □ □□□ □□ □□□□ □□□□ □□ □□.

□□□ □□ □□□ □□□□ □□□□ □□□□□?

A. SQS □□□□ □□□□ Lambda □□□ □□□□ □□□ □□□□ □□□□□ AWS Lambda□ □□□□□□. □□□ □□□□ □□□□ □□□□ □□□□ □□□ □ □□□ □□□□ □□ □□□ □□□ Amazon S3 □□□□ □□□□. □□□□ □□□□ □□□ □ □□□ □□ Lambda □□□ □□□□ SQS □□□□ □□□□ □□□□□.

B. SQS □□ □□□□ SQS FIFO □□□□ □□□□□□. Amazon EventBridge □□□ □□□□ 10□□□□ SQS □□□□ □□□□□□ AWS Lambda□ □□□□□□. Lambda □□□ □□□□ 5□□□ □□□ SentTimestamp □□ □□ □□□□ □□□□ □□□□ □□□□□□□□ □□ □□□ □□□ □□□ □□□□ □□□□□ □□□□ □□□□□.

C. SQS □□ □□ □□ □□□□ □□□□□□. Maximum Receives □□□ 1□ □□□□ □□ □□ □□ □□□ ARN□ □□ □□□ □□□ □ ARN□□ □□□□ □□□□□□ □□□ □□□□ □□ □□□□ □□□□□□. □□□□ □□ □□□□ □□□□ □□ □□ □□ □□ □□ □□ □□□□□ □□□□□□ □□□□□□. □□□ □ □□□□ □□ □□□□□□.

D. □ □□□ □□ □□□ □□□ □□ SQS □□ □□□□ □□□□ □□□□□□ API Gateway□ □□□□□□. □□□ □ □□ □□ □□□□ □□ □ □□ □□ □□□□□ □□□□□□□□ □□□□□□□□ □□□□ □ □□ □□ □□□□. □□□□□□□ □□ □□□□ □□□□ □□ □□ □□□□ □□□□□ □□□□□□. □□□ □ □□□□ □□ □□□□□□□.

Answer: C (LEAVE A REPLY)

Create an SQS dead-letter queue. Modify the existing queue by including a redrive policy that sets the Maximum Receives setting to 1 and sets the dead-letter queue ARN to the ARN of the newly created queue.

Instruct the scientists to use the dead-letter queue to review the data that is not valid. Reprocess this data at a later time.

NEW QUESTION: 5

□□□□□ □□ AWS CodePipeline□ □□□□□ Java □□□□□□□ □□□ □□□□□□□ □□□□□ □□□□□. □□□□□□□ □□ □□, □□ □□, □□ □□□ □□□□□□. □ □□□□□ runOrder □□ 1□ □□ □□□ □□□□□ □□□□□.

□□ □□ □□□□ □□ □□□ □□□□□□□ □□□□□ □□□. □□□□ □□ □□ □□□□ □□□ □□ □□ □□□ □□□□ □□□□ □□□□□.

□□□ □□ □□□ □□□□ □□□□ □□□□□?

A. □□ □□□ □□□□□□. runOrder □□ 1□ □□□ □□□ □□□□□□. AWS CodeDeploy□ □□ □□□□□ □□□□ □□ □□□□ □ □□□□□.

B. □□ □□ □□ runOrder □□ 2□ □□□ □□ □□ AWS CodeBuild□ □□ □□□□□ □□□□□ □□ □□□ □□

C. □□ □□ □□ runOrder □□ 1□ □□□ □□ □□ AWS CodeDeploy□ □□ □□□□□ □□□□□ □□ □□□□ □□

D. runOrder 2 AWS CodeBuild

Answer: B (LEAVE A REPLY)

Modify the Build Stage to Add a Test Action with a RunOrder Value of 2:

The build stage in AWS CodePipeline can have multiple actions. By adding a test action with a runOrder value of 2, the test action will execute after the initial build action completes.

Use AWS CodeBuild as the Action Provider to Run Unit Tests:

AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages.

Using CodeBuild to run unit tests ensures that the tests are executed in a controlled environment and that only the code changes that pass the unit tests proceed to the deploy stage.

Example configuration in CodePipeline:

```
{
  "name": "BuildStage",
  "actions": [
    {
      "name": "Build",
      "actionTypeId": {
        "category": "Build",
        "owner": "AWS",
        "provider": "CodeBuild",
        "version": "1"
      },
      "runOrder": 1
    },
    {
      "name": "Test",
      "actionTypeId": {
        "category": "Test",
        "owner": "AWS",
        "provider": "CodeBuild",
        "version": "1"
      },
      "runOrder": 2
    }
  ]
}
```

By integrating the unit tests into the build stage and ensuring they run after the build process, the pipeline guarantees that only code changes passing all unit tests are deployed.

References:

AWS CodePipeline

AWS CodeBuild

Using CodeBuild with CodePipeline

NEW QUESTION: 6

DevOps wants to use AWS IAM Identity Center (AWS Single Sign-On) to connect to an on-premises Active Directory (AD) instance. The AD instance uses SAML 2.0 for authentication.

DevOps wants to use AWS IAM Identity Center (AWS Single Sign-On) to connect to an on-premises Active Directory (AD) instance. The AD instance uses SAML 2.0 for authentication.

DevOps wants to use AWS IAM Identity Center (AWS Single Sign-On) to connect to an on-premises Active Directory (AD) instance. The AD instance uses SAML 2.0 for authentication.

- A. IAM Identity Center uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.
- B. IAM Identity Center uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.
- C. IAM Identity Center uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.
- D. IAM Identity Center uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.
- E. IAM Identity Center uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.
- F. IAM Identity Center uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.

Answer: (SHOW ANSWER)

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 `PrincipalTag`. 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: 7

DevOps wants to use AWS CodePipeline to build and deploy applications. The build step uses AWS CodeBuild. The build artifacts are stored in Amazon S3. The build artifacts are expired after 90 days.

DevOps wants to use AWS CodePipeline to build and deploy applications. The build step uses AWS CodeBuild. The build artifacts are stored in Amazon S3. The build artifacts are expired after 90 days.

- A. CodeBuild uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.
- B. CodeBuild uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.
- C. CodePipeline uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.
- D. CodeBuild uses the `aws:PrincipalTag` attribute to map AD groups to IAM roles.

Answer: B (LEAVE A REPLY)

The correct solution is to add a report group in the AWS CodeBuild project `buildspec` file with the appropriate path and format for the reports. Then, create an Amazon S3 bucket to store the reports. You should configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is completed. Finally, create an S3 Lifecycle rule to expire the objects

after 90 days. This approach allows for the automated transfer of reports to long-term storage and ensures they are retained for the required duration without manual intervention¹.

AWS CodeBuild User Guide on test reporting¹.

AWS CodeBuild User Guide on working with report groups².

AWS Documentation on using AWS CodePipeline with AWS CodeBuild³.

NEW QUESTION: 8

A SaaS application uses Amazon Elastic Load Balancing (ALB) to route traffic to Amazon Elastic Container Service (Amazon ECS) tasks. The application is deployed using AWS CodePipeline and AWS CodeDeploy. The application is currently running on a single Amazon EC2 instance. The application is deployed to a new Amazon EC2 instance. Which configuration should be used to ensure that traffic is shifted automatically in equal increments over a defined total deployment time with no manual intervention?

A. CodeDeploy with the `AppSpec` configuration `appspectra.yml` and the `TrafficRoutingConfig` configuration `TimeBasedLinear` with the `linearPercentage` and `linearInterval` options.

B. CodeDeploy with the `AppSpec` configuration `appspectra.yml` and the `TrafficRoutingConfig` configuration `AllAtOnce`.

C. CodeDeploy with the `AppSpec` configuration `appspectra.yml` and the `TrafficRoutingConfig` configuration `TimeBasedCanary` with the `initialPercentage` and `interval` options.

D. CodeDeploy with the `AppSpec` configuration `appspectra.yml` and the `TrafficRoutingConfig` configuration `WeightedRoundRobin`.

Answer: A (LEAVE A REPLY)

Comprehensive and Detailed Explanation From Exact Extract of DevOps Engineer documents only:

The requirement is to shift traffic automatically in equal increments over a defined total deployment time with no manual intervention. In AWS CodeDeploy blue/green deployments for Amazon ECS, the `TimeBasedLinear` traffic routing option is specifically designed for this purpose.

Option A is correct because:

- * `TimeBasedLinear` shifts traffic in equal percentages.
- * The deployment proceeds automatically at each interval.
- * `linearPercentage` defines how much traffic is shifted each time.
- * `linearInterval` defines how often the traffic shift happens.
- * This supports gradual traffic shifting and zero-downtime blue/green deployment behavior.

Why the other options are incorrect:

B). `AllAtOnce` shifts all traffic immediately, not in equal increments over time.

C). `TimeBasedCanary` shifts traffic in an initial portion and then the remainder later. That is not equal incremental shifting across the full deployment period.

D). This option is too generic and does not specifically identify the CodeDeploy configuration that provides equal incremental automatic traffic shifting. The exact feature required is `TimeBasedLinear`.

NEW QUESTION: 9

A company is using AWS CloudFront to deliver content to users. The content is stored in Amazon S3. The company wants to ensure that the content is always available and that the delivery is fast. Which of the following is the best way to ensure that the content is always available and that the delivery is fast?

A. Amazon CodeGuru is a code quality service that helps you find errors in your code. It uses static analysis to identify errors in your code. It is available for Python and Java. It is a managed service that you can use to improve the quality of your code. It is available in the AWS Systems Manager Parameter Store. SAM is a framework for deploying serverless applications. It is available in the AWS Systems Manager Parameter Store.

B. CodeCommit is a managed source control service that helps you manage your code. It is available in the AWS Systems Manager Parameter Store. SAM is a framework for deploying serverless applications. It is available in the AWS Systems Manager Parameter Store.

C. Amazon CodeGuru is a code quality service that helps you find errors in your code. It uses static analysis to identify errors in your code. It is available for Python and Java. It is a managed service that you can use to improve the quality of your code. It is available in the AWS Systems Manager Parameter Store. SAM is a framework for deploying serverless applications. It is available in the AWS Systems Manager Parameter Store.

D. CodeCommit is a managed source control service that helps you manage your code. It is available in the AWS Systems Manager Parameter Store. SAM is a framework for deploying serverless applications. It is available in the AWS Systems Manager Parameter Store.

Answer: (SHOW ANSWER)

<https://docs.aws.amazon.com/codecommit/latest/userguide/how-to-amazon-codeguru-reviewer.html>

NEW QUESTION: 11

You are using AWS Organizations to manage your AWS accounts. You want to ensure that all IAM users in all accounts have a policy that denies them access to the root user. Which of the following policies would you use to accomplish this goal?

A. A policy that denies access to the root user in all accounts.

B. A policy that denies access to the root user in all accounts, except for the root user in the master account.

C. A policy that denies access to the root user in all accounts, except for the root user in the master account and the root user in the member accounts.

D. A policy that denies access to the root user in all accounts, except for the root user in the master account and the root user in the member accounts, and the root user in the member accounts.

A.

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


B. CloudFront □□□ □ ALB□ □□□□□ □□□□□. □□□ □□ □□ □□□ □□□□□. AMI□ DR □□□ □□□□□. DR □□□ □□ □□□□ □□□ □□□ 0□ □□ □□ □□□ □□□□□. DR □□□ □ OpenSearch □□□ □□□□□ □□□□□. □□□□□ □ □□□□ □ □□□ □□□□□.

C. DR □□□ □ CloudFront □□□ □□□□ □ ALB□ □□□□□ □□□□□. □□ □□ □□(failover)□□ Amazon Route 53 DNS□ □ □□□□. AMI□ DR □□□ □□□□□. DR □□□ □□ □□□□ □□□ □□□ 0□ Auto Scaling □□□ □□□□□. OpenSearch □□ □ □□□□□ □□ AZ(Standby □□ □□)□ □□□□□□□□. □□ □□□ DR □□□ □□□ □□□□□.

D. DR □□□ □ CloudFront □□□ □□□□ □ ALB□ □□□□□ □□□□□. □□ □□ □□(failover)□□ Amazon Route 53 DNS□ □ □□□□. AMI□ DR □□□ □□□□□. DR □□□ □□ □□□□ □□□ □□(3)□ □□ Auto Scaling □□□ □□□□□. OpenSearch □□□ □□□□□ □□ □□□ □□□□ □□ AZ□ □□□□□□□□. □□ □□□ DR □□□ □□□ □□□□□.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 13

□ □□□ □□□ □□ □□□□□□ □□□ □□□ □□□□ □ □□□□□□□ □□□□ □□□□. □ □□□ □□□ □□ □□ □□□ □ □□□□. □□□ JSON □□□□□□. □ □□□□□□□ Amazon CloudWatch Logs □□ □□□ □□□ □□□□□□.

□□□ □□□□ □□ □□□ □□□□ □□□□ □□□ □□□□□.

A. Amazon OpenSearch □□□ □□□□□ OpenSearch □□□ □□ □□□ □□□□ □□ □□ □□□□ □□□□□ □□□□□. OpenSearch □□□ □□□□□ □□□□ □□ □□ □□ □□□ □□□□□ □□ □□ □□□ □□□□□ □□□□□.

B. □□□□ □□ □□□ □□□□ □□ □□□ □□ □□ □□□□ □□ □□□ □□ CloudWatch □□□ □□□ □□□□□. □ □□□□ □□□□□ □□□ □□□□ □□□□□ CloudWatch □□□□□ □□□□□.

C. AWS Lambda □□□ □□□□□ □□ □□□ □□ CloudWatch □□ □□□ □□□□.

Lambda □□□ □□□□□ JSON □□□ □□ □□□□ □□ □□□ □□□□□ □□ □□□ □□ □□□ CloudWatch□ □□□□□. □□ □ □□ □□□ □□ □□□□ □□□□□ CloudWatch □□□□□ □□□□□.

D. □□ □□□ □□□□□ Amazon Kinesis □□□ □□□□□ □□□□□. □□ □□□ □□□□ □□ □□ □□□□ □□□□□ □□□□ □□□ □□□ □□□□□ □□□ □□□□□ □□□□□. Lambda □□□ □□□ □□ □□□ CloudWatch□ □□□□□ □□ □□□□. □□□ □□ □□□ □□□□□ □□ □□□□ □□□□□ CloudWatch □□□□□ □□□□□.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 14

□ □□□□ □□ □□□□□□□□ AWS CodeDeploy □□□ □□□ □□□□□□. □□ □□□□□□□ Amazon EC2 □□□□□ □□□ □□□□□. DevOps □□□□□ □□□ □□□ □□□□ □□□ □□ □□□ □□□□ □□□.

□□ □□□□ □□ □□□ □□□□ □ □□□ □□□ □□□ □□□□□?

A. □□□□ □□□□ □□□□□□□ VPC □□ □□□ □□□□□. Amazon Inspector□ □□□□ □□□□ □□□ □□□□ □□ □□□ □□□□□. Amazon Detective□ □□□□ □□□ □□□□□.

B. EC2 □□□□□ □□ □□ □□□□□ □□□□□□. AWS Systems Manager Run Command□ □□□□ □□ EC2 □□□□□□ □□ □□ □□ □□□□□ □□□□□. AWS CloudTrail □□□□□ □□□ □□□□□.

C. Amazon CloudWatch Logs□ □□□□ □□□□□□ □□□ □□□□□. EC2 □□□□□ /opt/codedeploy-agent/deployment-root/ □□ □□□□ CodeDeploy □□ □□□ □□□□□. AWS X-Ray□ □□□□ □□□□□□□ □□ □□□ □□□ □□□□□.

D. CodeDeploy □□□ □□ AWS Trusted Advisor □□ □□□ □□□□□. AWS Health Dashboard□ □□□□ □□□□□□ □□□ □□ □□□□□. Amazon CloudWatch □□□□□□ □□ □□□ □□□□□.

Answer: C (LEAVE A REPLY)

CodeDeploy deployment failures must be diagnosed using deployment-specific logs and related application logs. CodeDeploy writes detailed logs on the target EC2 instances under /opt/codedeploy-agent/deployment-root/ (for example in logs and deployment-logs subdirectories). These logs include information about lifecycle events (BeforeInstall, AfterInstall, ApplicationStart, etc.), script execution failures, permissions issues, and any exit codes returned by deployment hooks.

Option C correctly points to CodeDeploy agent logs on the EC2 instances as the primary source for root cause analysis. Additionally, application-level logs streamed to Amazon CloudWatch Logs help validate whether the application itself is working correctly after deployment. For complex, multi-tier applications, AWS X-Ray can trace requests end to end, helping determine if issues are caused by downstream dependencies instead of the deployment process.

Options A, B, and D focus on network flow, generic performance, or high-level checks, none of which provide sufficient, deployment-level detail for CodeDeploy. Trusted Advisor and AWS Health do not surface step-by-step deployment log data. Therefore, the combination of CodeDeploy logs on the instances, CloudWatch Logs, and (optionally) X-Ray is the correct troubleshooting approach.

NEW QUESTION: 15

A DevOps team is migrating a multi-tier application from on-premises to AWS. The application consists of an Amazon EC2 instance that connects to an Amazon S3 bucket and an Amazon Lambda function. The team wants to ensure that the application can be deployed and managed consistently across different environments. Which of the following actions should the team take to achieve this goal?

- A. Configure AWS Systems Manager on each EC2 instance to manage the application. Configure AWS Lambda to connect to the Amazon S3 bucket. Configure AWS Config to monitor the application.
- B. Configure AWS Systems Manager on each EC2 instance to manage the application. Configure AWS Config to monitor the application.
- C. Configure AWS Systems Manager on each EC2 instance to manage the application. Configure AWS CloudTrail to monitor the application.
- D. Amazon CloudWatch Logs can be used to monitor the application. Amazon CloudWatch Logs can be used to monitor the application. Amazon CloudWatch Logs can be used to monitor the application.

Answer: A (LEAVE A REPLY)

Configure AWS Systems Manager on Each Instance:

AWS Systems Manager provides a unified interface for managing AWS resources. Install the Systems Manager agent on each EC2 instance to enable inventory management and other features.

Use AWS Systems Manager Inventory:

Systems Manager Inventory collects metadata about your instances and the software installed on them. This data includes information about applications, network configurations, and more.

Enable Systems Manager Inventory on all EC2 instances to gather detailed information about installed applications.

Use Systems Manager Resource Data Sync to Synchronize and Store Findings in an Amazon S3 Bucket:

Resource Data Sync aggregates inventory data from multiple accounts and regions into a single S3 bucket, making it easier to query and analyze the data.

Configure Resource Data Sync to automatically transfer inventory data to an S3 bucket for centralized storage.

Create an AWS Lambda Function that Runs When New Objects are Added to the S3 Bucket:

Use an S3 event to trigger a Lambda function whenever new inventory data is added to the S3 bucket.

The Lambda function can parse the inventory data and check for the presence of prohibited applications.

Configure the Lambda Function to Identify Prohibited Applications:

The Lambda function should be programmed to scan the inventory data for any known prohibited applications and generate alerts or take appropriate actions if such applications are found.

Example Lambda function in Python

```
import json
```

```
import boto3
```

```
def lambda_handler(event, context):
```

```
    s3 = boto3.client('s3')
```

```
    bucket = event['Records'][0]['s3']['bucket']['name']
```

```
    key = event['Records'][0]['s3']['object']['key']
```

```
    response = s3.get_object(Bucket=bucket, Key=key)
```

```
    inventory_data = json.loads(response['Body'].read().decode('utf-8')) prohibited_apps = ['app1', 'app2'] for instance in
```

```
    inventory_data['Instances']:
```

```
    for app in instance['Applications']:
```

```
    if app['Name'] in prohibited_apps:
```

```
    # Send notification or take action
```

```
    print(f" Prohibited application found: {app['Name']} on instance {instance['InstanceId']} " ) return { 'statusCode' : 200, 'body' :
```

```
    json.dumps( ' Check completed ' )} By leveraging AWS Systems Manager Inventory, Resource Data Sync, and Lambda, this solution
```

```
provides an efficient and automated way to audit EC2 instances for prohibited applications.
```

References:

AWS Systems Manager Inventory

AWS Systems Manager Resource Data Sync

S3 Event Notifications

AWS Lambda

NEW QUESTION: 16

□□□ Auto Scaling □□□ □□ Amazon EC2 □□□□ □□ Application Load Balancer□ □□□□ □□□□□□□ □□□□□. □□□□ □□□□ □□ □□□□□□. Auto Scaling □□□ □□□ □□ □□□ □□□ □□ AMI□ □□□□□. EC2 □□□□□□ □□□ □□ □□□□□ □□□□□ □□□□□.

Auto Scaling □□□□ □□□□ AMI□ □□ □□□□□□□. AMI ID□ □□□□ □□ □□□ Auto Scaling □□□ □□ □□□ □□□ □□□□.

DevOps □□□□□ □□□ □□ □□□ □□□□□ □□ □□ □□□ □□□□ □□□? (3□□ □□□□□.)

A. □ AMI□ □□□□ □ □□ □□□□ □□□□□.

B. □ □□ □□□□ □□□□□ Auto Scaling □□□ □□□□□□□.

C. Auto Scaling □□□ □□□ □□□ □□ □□□□.

D. Auto Scaling □□□ □□□ □□□ I□□ □□□□.

E. Auto Scaling □□□□ □□ □□ EC2 □□□□□□ □ AMI□ □□□□□.

F. EC2 □□□□□ □□□□ □□ □□□ □□ □□□ AMI□ □□□□ □ AMI□ □□□□□.

Answer: A,B,F ([LEAVE A REPLY](#))

- □□ □□□ □□□□ □□ DevOps □□□□□ □□□□ □□ □□ □□□ □□□□□? (□ □□□ □□□□□.)
- A. AWS Config □□ □□□ □□□ □□□ □□□□□. □□□□ AWS Config □□ □□□ □ □□ □□□□ □□□□□□.
 - B. AWS Config □□ □□□ □□□ □□□ □□□□□. □□□ □□□ □□□□ AWS Config □□ □□□ □□ □□□ □□□□□.
 - C. AWS CloudFormation □□□□ □□□□ AWS Config □□□□ □□□□□. □□□ □□ □□□ □□□□ □□□□□ CloudFormation □□ □□□ □□□□□.
 - D. □□□ □□□ □□□□ AWS Config □□ □□□□ □□□□□. □□□ □□ AWS □□ □ □□ AWS □□□□ □□□ □□□ □□□□ □.
 - E. □□□ □□□ □□□□ AWS Config □□ □□□□ □□□□□. □□□ □□ AWS □□ □ □□ AWS □□□□ □□□ □□□ □□□□ □.

Answer: A,E (LEAVE A REPLY)

<https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/> <https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html>

NEW QUESTION: 19

- □□□ □□ AWS □□□ □□□□ □□ □□ □□□ □□□□ AWS Organizations □□□ □□□□ □□□□. □ □□□ □□ □□□ AWS Config □□□□□□□□. □□□□□□ □□ □□ □□ □□□□□□□□ □□□ □□□□□ □□ □□□ AWS □□□ AWS CloudFormation □□□ □□□□ □□□.
- □□□ □□□□□ □□□□□□□□ □□□ Amazon EC2 □□□□ □□□ □□□□□ □□□□□ □□□.
- □□□□ □□□ □□ □□□ □□□□□□?
- A. EC2 □□□□ □□ □□□ □□□ □□□□ □□ □□□ □□ □□□ □ SUCCESS □ □□□□ AWS Lambda □□□ □□□□□. □ AWS □□□□ CloudFormation Guard Hook □ □□□□□□□ Lambda □□□ □□□□□.
 - B. □ AWS □□□□ desired-instance-type □□□ □□□□ AWS Config □□□ □□□□□. □□ □□□ □□□ □□□□ □□ □□□ □□□□. AWS-StopEC2Instance □□ □□□ □□□□ AWS Config □□□ □□ □□ □□□ □□□□□.
 - C. ec2:InstanceType □□□ □□□ □□□□ □□ □□□ □□ □□□□ □□ □ ec2:RunInstances □ □□ □□ □□□ □□□□ SCP □ □□□□□. □□ SCP □ □□□ □□□ □□□□□.
 - D. EC2 □□□□ □□□ □□□ □□□□ □□ □□□ □□ □□□□□ □□□□ CloudFormation Guard □□□ □□□□□. □ AWS □□ □□ Guard Hook □ □□□□□□□ Guard □□□ □□□□□.

Answer: D (LEAVE A REPLY)

The key requirement is to prevent non-approved EC2 instance types from being created at stack creation time, specifically when developers deploy AWS CloudFormation stacks. This is a pre-deployment enforcement requirement, not a detection or remediation requirement after resources already exist.

CloudFormation Guard Hooks are purpose-built for this use case. They allow organizations to define policy- as-code rules that validate CloudFormation templates before resources are created or updated. By writing a CloudFormation Guard rule that explicitly checks the InstanceType property against an approved allow list, the stack operation will fail immediately if a developer attempts to use a disallowed instance type. This enforces compliance early in the deployment lifecycle and avoids post-deployment cleanup.

Option B uses AWS Config, which only detects noncompliance after resources are created. Even with remediation, the instance would still be launched briefly, which violates the requirement. Option C uses an SCP, which applies broadly to all EC2 launches in the organization and is not limited to CloudFormation stacks; this is overly restrictive and could unintentionally block other valid use cases. Option A incorrectly combines Lambda with Guard Hooks-Guard Hooks natively evaluate Guard rules and do not invoke Lambda functions. Therefore, using a CloudFormation Guard rule with a Guard Hook is the correct and AWS-recommended solution for enforcing approved EC2 instance types during CloudFormation deployments.

NEW QUESTION: 20

A company is using Amazon S3 to store log files. The logs are generated by an application that runs on Amazon EC2 instances. The logs are stored in a bucket and are accessed by a web application. The web application needs to access the logs and the logs are sensitive. The company wants to ensure that the logs are protected and that the web application can only access the logs it needs. What is the best way to protect the logs and control access to them?

- A. Create an IAM role for the web application and grant it access to the logs using S3 access points. Create a separate IAM role for the EC2 instances and grant it access to the logs using S3 access points. Use S3 access points to control access to the logs.
- B. Amazon Kinesis can be used to stream the logs to a Kinesis Data Store. The Kinesis Data Store can be accessed by the web application using AWS Lambda. The logs can be processed by the Lambda function and the results can be stored in a separate bucket.
- C. The logs can be encrypted using S3 Object Lambda. The web application can access the logs using S3 Object Lambda. The logs can be processed by the Lambda function and the results can be stored in a separate bucket.
- D. The logs can be encrypted using S3 Object Lambda. The web application can access the logs using S3 Object Lambda. The logs can be processed by the Lambda function and the results can be stored in a separate bucket.

Answer: D (LEAVE A REPLY)

The best solution is to use S3 Object Lambda¹, which allows you to add your own code to S3 GET, LIST, and HEAD requests to modify and process data as it is returned to an application². This way, you can redact the data differently for each application without creating and storing multiple copies of the data or running proxies.

The other solutions are less efficient or scalable because they require replicating the data to multiple buckets, streaming the data through Kinesis, or storing the data in S3 access points.

1: Amazon S3 Features | Object Lambda | AWS 2: Transforming objects with S3 Object Lambda - Amazon Simple Storage Service

NEW QUESTION: 21

A company is using AWS Organizations to manage multiple AWS accounts. The company wants to ensure that the accounts are managed in a secure and compliant way. The company also wants to ensure that the accounts are managed in a way that is easy to set up and govern. What is the best way to manage the accounts?

- A. AWS Config can be used to manage the accounts. The company can use the account-part-of-organizations parameter in the AWS Config console to manage the accounts.
- B. Enterprise Support can be used to manage the accounts. The company can use the support:ResolveCase parameter in the AWS Lambda console to manage the accounts.
- C. control_tower_parameters can be used to manage the accounts. The company can use the AWSEnterpriseSupport parameter in the AWS Lambda console to manage the accounts.
- D. AFT can be used to manage the accounts. The company can use the aft_feature_enterprise_support parameter in the AWS Lambda console to manage the accounts.

Answer: (SHOW ANSWER)

AWS Organizations is a service that helps to manage multiple AWS accounts. AWS Control Tower is a service that makes it easy to set up and govern secure, compliant multi-account AWS environments. Account Factory for Terraform (AFT) is an AWS Control Tower feature that provisions new accounts using Terraform templates. To provision new accounts with the Enterprise Support plan, the DevOps

engineer can set the `aft_feature_enterprise_support` feature flag to True in the AFT deployment input configuration. This flag enables the Enterprise Support plan for newly provisioned accounts.

<https://docs.aws.amazon.com/controltower/latest/userguide/aft-feature-options.html>

NEW QUESTION: 22

A DevOps engineer is configuring Amazon Macie for an AWS account. The engineer wants to optimize Macie costs for the account without compromising the account's functionality. The 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.

- A. 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.
- B. 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.
- C. 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.
- D. 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.
- E. 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.

Answer: (SHOW ANSWER)

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: 23

A DevOps engineer is configuring Amazon GuardDuty for an AWS account. The engineer wants to immediately detect and terminate EC2 instances involved in cryptocurrency mining with the least development effort. Amazon GuardDuty is the AWS-native service specifically designed to detect malicious activities such as crypto-mining by continuously analyzing CloudTrail events, VPC Flow Logs, and DNS logs. GuardDuty includes managed threat intelligence and predefined findings like `CryptoCurrency:EC2/BitcoinTool.B!DNS` and `CryptoCurrency:EC2/BitcoinTool.B!IP`, which directly identify mining behavior without custom detection logic.

The requirement is to immediately detect and terminate EC2 instances involved in cryptocurrency mining with the least development effort. Amazon GuardDuty is the AWS-native service specifically designed to detect malicious activities such as crypto-mining by continuously analyzing CloudTrail events, VPC Flow Logs, and DNS logs. GuardDuty includes managed threat intelligence and predefined findings like `CryptoCurrency:EC2/BitcoinTool.B!DNS` and `CryptoCurrency:EC2/BitcoinTool.B!IP`, which directly identify mining behavior without custom detection logic.

- A. Amazon Route 53 is the AWS-native service specifically designed to detect malicious activities such as crypto-mining by continuously analyzing CloudTrail events, VPC Flow Logs, and DNS logs. GuardDuty includes managed threat intelligence and predefined findings like `CryptoCurrency:EC2/BitcoinTool.B!DNS` and `CryptoCurrency:EC2/BitcoinTool.B!IP`, which directly identify mining behavior without custom detection logic.
- B. VPC is the AWS-native service specifically designed to detect malicious activities such as crypto-mining by continuously analyzing CloudTrail events, VPC Flow Logs, and DNS logs. GuardDuty includes managed threat intelligence and predefined findings like `CryptoCurrency:EC2/BitcoinTool.B!DNS` and `CryptoCurrency:EC2/BitcoinTool.B!IP`, which directly identify mining behavior without custom detection logic.
- C. Amazon GuardDuty is the AWS-native service specifically designed to detect malicious activities such as crypto-mining by continuously analyzing CloudTrail events, VPC Flow Logs, and DNS logs. GuardDuty includes managed threat intelligence and predefined findings like `CryptoCurrency:EC2/BitcoinTool.B!DNS` and `CryptoCurrency:EC2/BitcoinTool.B!IP`, which directly identify mining behavior without custom detection logic.
- D. AWS Security Hub is the AWS-native service specifically designed to detect malicious activities such as crypto-mining by continuously analyzing CloudTrail events, VPC Flow Logs, and DNS logs. GuardDuty includes managed threat intelligence and predefined findings like `CryptoCurrency:EC2/BitcoinTool.B!DNS` and `CryptoCurrency:EC2/BitcoinTool.B!IP`, which directly identify mining behavior without custom detection logic.

Answer: C (LEAVE A REPLY)

The requirement is to immediately detect and terminate EC2 instances involved in cryptocurrency mining with the least development effort. Amazon GuardDuty is the AWS-native service specifically designed to detect malicious activities such as crypto-mining by continuously analyzing CloudTrail events, VPC Flow Logs, and DNS logs. GuardDuty includes managed threat intelligence and predefined findings like `CryptoCurrency:EC2/BitcoinTool.B!DNS` and `CryptoCurrency:EC2/BitcoinTool.B!IP`, which directly identify mining behavior without custom detection logic.

Option C leverages this built-in capability. Once GuardDuty is enabled, findings are automatically generated when mining activity is detected. These findings are sent to Amazon EventBridge in near real time. An EventBridge rule can filter for cryptocurrency-related findings and trigger an AWS Lambda function . The Lambda function can then identify the affected EC2 instance and terminate it or adjust the Auto Scaling group to replace it. This approach requires minimal custom code and no log parsing, scheduled jobs, or analytics pipelines.

Options A and B rely on custom log analysis, periodic execution, and maintaining lists of mining domains or IPs, which significantly increases complexity and response time. Option D uses AWS Security Hub, which aggregates findings from GuardDuty and other services but is not intended for immediate, low-latency remediation.

Therefore, Option C provides the fastest detection, immediate response, and lowest development overhead using AWS-managed threat detection services.

NEW QUESTION: 24

□□ □□ AWS CodeCommit □ □□□□ □□□□□□ □□□ □□□ □□□□ AWS CodePipeline □ □□□□ □□□□□ □□□ □□□ □ □□□□. □□ □□ □□ □□□ □□□□ □□□□□□ □□□□ □□ □□ □□□ □□□□□ □□□□□□. □□□□□ CodeCommit □□□□□□ □□ □□ □□□ □□□□□ □□□□□□ 10□ □□□ □□□ □□□ □□□□□□.

□ □□□ □□□□□ □□ □ □□ □□□ □□□ □□□?

- A. □□ □□□□ □□□□□□ □□□□□□ Amazon EventBridge □□□ □□□□□□ □□□□□□.
- B. CodePipeline □□□ □□□ CodeCommit □□□□□□ □□□□ □ □□ □□□ □□□ □□□□□□.
- C. □□□□ IAM □□□ CodeCommit □□□□□□ □□□ □ □□ □□□ □□□ □□□□□□.
- D. Amazon CloudWatch Logs □ CodeCommit □□□ □□ □□□□□□ □□□□ □□□□ □□□□□.

Answer: (SHOW ANSWER)

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: 25

DevOps wants to use AWS CloudFormation to create an Amazon AD Connector. The CloudFormation template includes the `CREATE_IN_PROGRESS` and `CREATE_COMPLETE` hooks. What is the correct way to configure the hooks?

- A. Lambda function with IAM role and AWS Lambda function ds ConnectDirectory permissions.
- B. Lambda function with IAM role and AWS Lambda function ds ConnectDirectory permissions.
- C. Lambda function with IAM role and ARN of cloudformation UpdateStack permissions.
- D. Lambda function with URL of the CloudFormation console.

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 26

DevOps wants to use AWS CloudFormation to create an Amazon AD Connector. The CloudFormation template includes the `CREATE_IN_PROGRESS` and `CREATE_COMPLETE` hooks. What is the correct way to configure the hooks?

- A. CloudFormation StackSets with IAM role and AWS Lambda function ds ConnectDirectory permissions. EBS and SQS permissions are not required. CloudFormation Hook with StackSets OU and Hook permissions.
- B. OU with IAM role and AWS Config permissions. AWS Systems Manager OU and EBS and SQS permissions are not required. AWS Config permissions are not required.
- C. EBS and SQS permissions are not required. EBS and SQS permissions are not required. SCP OU and Hook permissions are not required.
- D. EBS and SQS permissions are not required. AWS Lambda function ds ConnectDirectory permissions are not required. Lambda function with IAM role and Hook permissions.

Answer: [A \(LEAVE A REPLY\)](#)

The requirement specifies enforcing encryption before CloudFormation creates or updates resources. This is key because preventive enforcement must occur during the provisioning workflow, not after resources already exist. AWS provides CloudFormation Hooks specifically for this purpose. A Hook allows an organization to intercept a CloudFormation stack operation and validate resource configurations before provisioning occurs.

This feature is recommended by AWS for pre-deployment governance such as enforcing encryption policies, tag compliance, or security restrictions.

By enabling trusted access between CloudFormation StackSets and AWS Organizations, the Hook can be deployed centrally from the delegated administrator account across all accounts in the specified OU. Any attempt to create or update EBS volumes or SQS queues through CloudFormation is validated first by the Hook. If encryption is not configured, the operation fails immediately.

Option C (SCP) blocks API calls globally, but SCPs cannot perform conditional logic based on resource properties passed by CloudFormation prior to creation. Option B (AWS Config) detects violations after resources already exist, which does not satisfy "before stack operation." Option D (Lambda remediation) also occurs after the resource is created.

Thus, CloudFormation Hooks distributed via StackSets provide the only solution that enforces compliance before the provisioning lifecycle

NEW QUESTION: 27

DevOps wants to monitor Amazon EKS pod-level metrics and logs. DevOps wants to monitor pod-level metrics and logs.

Which solution meets these requirements? DevOps wants to monitor pod-level metrics and logs.

Which solution meets these requirements?

A. EKS pod-level metrics and logs using CloudWatch Container Insights.

B. EKS pod-level metrics and logs using AWS CloudTrail.

C. EKS pod-level metrics and logs using CloudTrail Insights.

D. CloudWatch Observability add-on for EKS using Amazon CloudWatch Container Insights.

Answer: D (LEAVE A REPLY)

Container restart failures require detailed observability into pod-level, node-level, and container-level metrics and logs. CloudWatch Container Insights is purpose-built for Kubernetes operational diagnostics and provides granular visibility into CPU, memory, network I/O, disk I/O, container restarts, OOM kills, throttling, pod lifecycle issues, and Kubernetes control plane behaviors.

The CloudWatch Observability add-on deploys Fluent Bit and the CloudWatch Agent directly into the EKS cluster as DaemonSets. These components automatically collect:

Container logs

Pod metrics

Node metrics

Cluster events

OOM errors

CrashLoopBackOff restart cycles

Control plane request anomalies

With this data, the DevOps engineer can easily identify misconfigurations, resource bottlenecks, unhealthy nodes, failing containers, or image pull issues.

Option A (dashboards only) lacks per-container diagnostic data.

Option B (CloudTrail) only logs API calls - not useful for restart debugging.

Option C (CloudTrail Insights) only detects anomalous API usage, not container failures.

Therefore, CloudWatch Container Insights is the correct and AWS-recommended solution for diagnosing container restart failures in EKS.

NEW QUESTION: 28

DevOps wants to monitor AWS CloudTrail API calls. DevOps wants to monitor AWS CloudTrail API calls.

Which solution meets these requirements?

Action: Attach the CloudWatchAgentServerPolicy managed IAM policy to the IAM instance profile that the EKS cluster uses.

Why: This ensures the CloudWatch agent has the necessary permissions to collect memory metrics.

Reference: AWS documentation on CloudWatch Agent Permissions.

This corresponds to Option A: Attach the CloudWatchAgentServerPolicy managed IAM policy to the IAM instance profile that the cluster uses.

Step 2: Deploying the CloudWatch Agent to EC2 Instances To collect memory metrics from the EC2 instances running in the EKS cluster, the CloudWatch agent needs to be deployed on these instances. The agent collects system-level metrics, including memory usage.

Action: Deploy the unified Amazon CloudWatch agent to the existing EC2 instances in the EKS cluster.

Update the Amazon Machine Image (AMI) for future instances to include the CloudWatch agent.

Why: The CloudWatch agent allows you to collect detailed memory metrics from the EC2 instances, which is not enabled by default.

Reference: AWS documentation on Installing and Configuring the CloudWatch Agent.

This corresponds to Option C: Collect performance metrics by deploying the unified Amazon CloudWatch agent to the existing EC2 instances in the cluster. Add the agent to the AMI for any new EC2 instances that are added to the cluster.

Step 3: Analyzing Memory Metrics Using Container Insights After collecting the memory metrics, you can analyze them using the pod_memory_utilization metric in Amazon CloudWatch Container Insights. This metric provides visibility into the memory usage of the containers (pods) in the EKS cluster.

Action: Analyze the pod_memory_utilization CloudWatch metric in the Container Insights namespace by using the Service dimension.

Why: This provides detailed insights into memory usage at the container level, which helps diagnose memory-related issues.

Reference: AWS documentation on CloudWatch Container Insights.

This corresponds to Option E: Analyze the pod_memory_utilization Amazon CloudWatch metric in the Container Insights namespace by using the Service dimension.

NEW QUESTION: 31

□□□ Amazon EC2 □□□□□□ □□□□□□□ □□□□□ □□ □□ AWS □□□ □□□□□. □□□ AWS □□□□ AWS Config □□ limited-ssh AWS Config □□ □□□ □□□□□□□.

□□□ □□□ □□ □□□ □□□ SSH □□□ □□□□ □□ □□ □□□□□ □□□ □□□□ □□□□ □□□□ □□□□ □□ □□□. □□□ □□ □□□□ □□□ □□ □□□ □□□ ID□ □□□□□ □□□.

DevOps □□□□□ □□□□ Amazon Simple Notification Service(Amazon SNS) □□□ □□□□ □□□ □□□ □□□□ □□□□□.

DevOps □□□□□ □□□ □□ □□□ □□□□ □□ □□□ □□□ □□ □□□?

A. limited-ssh □□□ □□ NON_COMPLIANT□ AWS Config □□ □□□ □□□□ Amazon EventBridge □□□ □□□□□. EventBridge □□□ □□ □□ □□□ □□ SNS □□□ □□□ □□□□□ EventBridge □□□ □□□□□.

B. limited-ssh □□□ □□ □□ □□ □□□ SNS □□□ □□□□ AWS Config□ □□□□□. □□□ NON_COMPLIANT □□□□ □□ □□□ □□□□□ □□□□ SNS □□□ □□ □□ □□□ □□□□□.

C. limited-ssh □□□ □□ NON_COMPLIANT□ AWS Config □□ □□□ □□□□ Amazon EventBridge □□ □□ SNS □□□□ AWS Systems Manager Run Command□ □□□□ □□□ □□□ □□□□ □□□ SNS □□

D. NON_COMPLIANT□ □□ AWS Config □□ □□□ □□□□ Amazon EventBridge □□□ □□□□□. limited-ssh □□□ □□ □□ □□□□ □□□□□. SNS □□□ □□□ □□□□□ EventBridge □□□ □□□□□.

Answer: A (LEAVE A REPLY)

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

B. `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

C. `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

" `aws:ResourceTag/access-team` " : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

D. `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

E. `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

F. `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

Answer: A,D,F (LEAVE A REPLY)

Short Explanation: To meet the requirements, the DevOps engineer should update the SAML assertion to pass the user's team name, update the IAM role's trust policy to add an access-team session tag that has the team name, create an IAM permissions boundary in each account, and for each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team. Updating the SAML assertion to pass the user's team name allows the DevOps engineer to use IAM tags to identify which team a user belongs to. This can help enforce fine-grained access control based on the user's team membership¹.

Updating the IAM role's trust policy to add an access-team session tag that has the team name allows the DevOps engineer to use IAM condition keys to restrict access based on the session tag value². For example, the DevOps engineer can use the `aws:PrincipalTag` condition key to match the access-team tag of the user with the access-team tag of the repository³.

Creating an IAM permissions boundary in each account allows the DevOps engineer to set the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries⁴. For example, the DevOps engineer can use a permissions boundary policy to limit the actions that a user can perform on CodeCommit repositories based on their access-team tag⁵.

For each CodeCommit repository, adding an access-team tag that has the value set to the name of the associated team allows the DevOps engineer to use resource tags to identify which team manages a repository. This can help enforce fine-grained access control based on the resource tag value⁶.

The other options are incorrect because:

Creating an approval rule template for each team in the Organizations management account is not a valid option, as approval rule templates are not supported by AWS Organizations. Approval rule templates are specific to CodeCommit and can only be associated with one or more repositories in the same AWS Region where they are created⁷.

Creating an approval rule template for each account is not a valid option, as approval rule templates are not designed to restrict access to modify branches. Approval rule templates are designed to require approvals from specified users or groups before merging pull requests⁸.

Attaching an SCP to the accounts is not a valid option, as SCPs are not designed to restrict access based on tags. SCPs are designed to restrict access based on service actions and resources across all users and roles in an organization's account⁹.

NEW QUESTION: 34

DevOps `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

`aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

A. `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

B. `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} " `aws:ResourceTag/access-team` : " \$;{aws:PrincipalTag/access-team} "

C. AWS Organizations SCPs that deny 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.

D. AWS Lambda 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.

E. 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.

Answer: A,B (LEAVE A REPLY)

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.

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: 35

DevOps team is migrating an application from on-premises to AWS. The application consists of a REST API, a database, and a set of microservices. The application is currently running on Amazon EC2 instances. The team wants to ensure that the application is highly available and resilient. They are considering using AWS CodeDeploy, AWS CodePipeline, AWS CodeBuild, and AWS Lambda. They are also considering using Amazon API Gateway, Amazon CloudWatch, and Amazon SNS. They are asking for your advice on how to best implement the application on AWS.

Which of the following options best meets the requirements?

A. CodeDeploy with LambdaAllAtOnce deployment configuration. API Gateway with Amazon CloudWatch alarms. SNS for notifications.

B. CodeDeploy with LambdaCanary10Percent10Minutes deployment configuration. API Gateway with Amazon CloudWatch alarms. SNS for notifications.

C. CodeDeploy with LambdaAllAtOnce deployment configuration. API Gateway with Amazon SNS Topic for notifications. SNS for notifications.

D. CodeDeploy with LambdaCanary10Percent10Minutes deployment configuration. API Gateway with Amazon CloudWatch alarms. SNS for notifications. SNS for notifications.

Answer: B (LEAVE A REPLY)

Option A is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring automatic rollbacks on the deployment group is not operationally efficient, as it requires manual intervention to fix the errors and redeploy the application.

Option B is correct because setting the deployment configuration to LambdaCanary10Percent10Minutes means that the new version of the application will be deployed to 10 percent of the Lambda functions first, and then to the remaining 90 percent after 10 minutes. This minimizes the impact of errors on customers, as only 10 percent of them will be affected by a faulty deployment. Configuring automatic rollbacks on the deployment group also meets the requirement of reverting to the most recent stable version of the application when an error is detected. Creating a CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway is a valid way to monitor the health of the application and trigger a rollback if needed.

Option C is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating an SNS topic to send notifications every time a deployment fails is not sufficient to detect errors in the application, as it does not monitor the API Gateway responses.

Option D is incorrect because configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating a metric filter on a CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors is a valid way to monitor the health of the application, but invoking a new Lambda function to perform a rollback is unnecessary and complex, as CodeDeploy already provides automatic rollback functionality.

References:

AWS CodeDeploy Deployment Configurations

[AWS CodeDeploy Rollbacks]

Amazon CloudWatch Alarms

NEW QUESTION: 36

You are implementing an organization trail with logs centralized in the audit account's S3 bucket. The trail must ensure compliance and isolation. An EventBridge rule in the audit account triggers on failed login events (ConsoleLogin failed) and sends SNS notifications in near real time.

Which of the following is the most efficient way to send the notification?

A. CloudTrail sends events to an S3 bucket. An EventBridge rule in the audit account triggers on failed login events (ConsoleLogin failed) and sends SNS notifications in near real time.

B. CloudTrail sends events to an Athena table. An EventBridge rule in the audit account triggers on failed login events (ConsoleLogin failed) and sends SNS notifications in near real time.

C. CloudTrail sends events to an S3 bucket. An EventBridge rule in the audit account triggers on failed login events (ConsoleLogin failed) and sends SNS notifications in near real time.

D. CloudTrail sends events to an S3 bucket. An EventBridge rule in the audit account triggers on failed login events (ConsoleLogin failed) and sends SNS notifications in near real time.

Answer: A (LEAVE A REPLY)

Using an organization trail with logs centralized in the audit account's S3 bucket ensures compliance and isolation. An EventBridge rule in the audit account triggers on failed login events (ConsoleLogin failed) and sends SNS notifications in near real time.

NEW QUESTION: 37

You are implementing an organization trail with logs centralized in the audit account's S3 bucket. The trail must ensure compliance and isolation. An EventBridge rule in the audit account triggers on failed login events (ConsoleLogin failed) and sends SNS notifications in near real time.

A. CloudTrail sends events to an S3 bucket. An EventBridge rule in the audit account triggers on failed login events (ConsoleLogin failed) and sends SNS notifications in near real time.

B. AWS CodeArtifact is a managed artifact repository that can be used to store and distribute artifacts. It is a fully managed service that integrates with AWS CodeBuild and AWS CodeDeploy. CodeArtifact is a fully managed service that integrates with AWS CodeBuild and AWS CodeDeploy.

C. AWS CodeDeploy is a managed service that automates the deployment of your applications. It integrates with Amazon Elastic File System (Amazon EFS) to store application artifacts.

D. AWS CodeBuild is a managed service that builds your source code. It integrates with AWS Artifact to store artifacts. CodeBuild is a managed service that builds your source code. It integrates with AWS Artifact to store artifacts.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 38

DevOps is a set of practices that enables an organization to release software updates more frequently and reliably. It integrates with Amazon Kinesis to stream data and AWS Lambda to process data. Lambda is a serverless compute service that runs code in response to events and can be triggered by REST APIs.

Lambda is a serverless compute service that runs code in response to events and can be triggered by REST APIs. It integrates with Amazon Simple Queue Service (Amazon SQS) to queue messages and DevOps to manage deployments. Lambda is a serverless compute service that runs code in response to events and can be triggered by REST APIs.

DevOps is a set of practices that enables an organization to release software updates more frequently and reliably. It integrates with REST APIs to manage deployments and Lambda to process data. DevOps is a set of practices that enables an organization to release software updates more frequently and reliably. It integrates with REST APIs to manage deployments and Lambda to process data.

Which of the following is the best solution to meet the requirements?

A. Use AWS Lambda to process data and REST APIs to manage deployments.

B. Use AWS Lambda to process data and Amazon SQS to queue messages.

C. Use AWS Lambda to process data and Amazon EFS to store artifacts.

D. Use AWS Lambda to process data and Amazon S3 to store artifacts.

Answer: B (LEAVE A REPLY)

This solution will meet the requirements because it will reduce the number of errorless records that are sent to the dead-letter queue.

When you configure the setting to split the batch when an error occurs, Lambda will retry only the records that caused the error, instead of retrying the entire batch. This way, the records that have no data errors and have already been processed by the legacy REST API will not be retried and sent to the dead-letter queue unnecessarily.

<https://docs.aws.amazon.com/lambda/latest/dg/with-kinesis.html>

NEW QUESTION: 39

Which of the following is the best solution to meet the requirements? The solution must be able to manage the configuration of AWS Organizations, AWS Inspector, AWS CloudFormation, and Amazon S3. The solution must be able to manage the configuration of AWS Organizations, AWS Inspector, AWS CloudFormation, and Amazon S3.

Which of the following is the best solution to meet the requirements?

A. Use AWS CloudFormation to manage the configuration of AWS Organizations, AWS Inspector, AWS CloudFormation, and Amazon S3. Use AWS CloudFormation to manage the configuration of AWS Organizations, AWS Inspector, AWS CloudFormation, and Amazon S3.

B. Use AWS Config to manage the configuration of AWS Organizations, AWS Inspector, AWS CloudFormation, and Amazon S3. Use AWS Config to manage the configuration of AWS Organizations, AWS Inspector, AWS CloudFormation, and Amazon S3.

C. Use AWS Config to manage the configuration of AWS Organizations, AWS Inspector, AWS CloudFormation, and Amazon S3. Use AWS Config to manage the configuration of AWS Organizations, AWS Inspector, AWS CloudFormation, and Amazon S3.

D. IAM roles for AWS Lambda and AWS CodePipeline. IAM roles for EBS and Amazon S3.

Answer: (SHOW ANSWER)

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

NEW QUESTION: 40

A DevOps engineer is designing a multi-region, multi-availability zone application. The application must be able to serve content from multiple Amazon CloudFront distributions. The application must also be able to serve content from an Amazon ElastiCache Redis instance. The application must be able to serve content from an Amazon ElastiCache Redis instance. The application must be able to serve content from an Amazon ElastiCache Redis instance.

The application must be able to serve content from an Amazon ElastiCache Redis instance. The application must be able to serve content from an Amazon ElastiCache Redis instance. The application must be able to serve content from an Amazon ElastiCache Redis instance. The application must be able to serve content from an Amazon ElastiCache Redis instance.

A. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region.

B. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region.

C. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region.

D. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region. Use Amazon ElastiCache Redis instances in each region.

Answer: B (LEAVE A REPLY)

The best solution to implement failover for the application is to use CloudFront origin groups. Origin groups allow CloudFront to automatically switch to a secondary origin when the primary origin is unavailable or returns specific HTTP status codes that indicate a failure¹. This way, CloudFront can serve the requests from the secondary ALB in the secondary Region without any delay or redirection. To set up origin groups, the DevOps engineer needs to create a new origin on the distribution for the secondary ALB, create a new origin group with the original ALB as the primary origin and the secondary ALB as the secondary origin, and configure the origin group to fail over for HTTP 5xx status codes. Then, the DevOps engineer needs to update the default behavior to use the origin group instead of the single origin².

The other options are not as effective or efficient as the solution in option B. Option A is not suitable because creating a second CloudFront distribution will increase the complexity and cost of the application. Moreover, using Route 53 alias records with a failover policy will introduce some delay in detecting and switching to the secondary CloudFront distribution, which may not meet the zero-second RTO requirement. Option C is not feasible because CloudFront does not support using Route 53 alias records as origins³. Option D is not advisable because using a CloudFront function to redirect the requests to the secondary ALB will add an extra round-trip and latency to the failover process, which may also not meet the zero-second RTO requirement.

1: Optimizing high availability with CloudFront origin failover - Amazon CloudFront

2: Creating an origin group - Amazon CloudFront

3: Values That You Specify When You Create or Update a Web Distribution - Amazon CloudFront

NEW QUESTION: 41

A DevOps engineer is designing a multi-region, multi-availability zone application. The application must be able to serve content from multiple Amazon CloudFront distributions. The application must also be able to serve content from an Amazon ElastiCache Redis instance. The application must be able to serve content from an Amazon ElastiCache Redis instance. The application must be able to serve content from an Amazon ElastiCache Redis instance.

Which of the following configurations meets the requirements?

A. AWS Serverless Application Model(AWS SAM) defines the Lambda function. AWS CodeDeploy deploys the function to the production environment. Amazon CloudWatch monitors the function. Amazon CloudWatch sends an alarm to the DevOps engineer when the function fails.

B. AWS CloudFormation defines the Lambda function. Amazon CloudWatch monitors the function. Amazon CloudFormation sends an alarm to the DevOps engineer when the function fails.

C. AWS CloudFormation defines the Lambda function. Amazon CloudWatch monitors the function. Amazon CloudWatch sends an alarm to the DevOps engineer when the function fails. The alarm is configured to send a message to the DevOps engineer when the function fails.

D. AWS CodeBuild builds the Lambda function. Amazon CloudWatch monitors the function. Amazon CloudWatch sends an alarm to the DevOps engineer when the function fails. The alarm is configured to send a message to the DevOps engineer when the function fails.

Answer: D (LEAVE A REPLY)

Use routing configuration on an alias to send a portion of traffic to a second function version. For example, you can reduce the risk of deploying a new version by configuring the alias to send most of the traffic to the existing version, and only a small percentage of traffic to the new version. [https://docs.aws.amazon.com](https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html)

[/lambda/latest/dg/configuration-aliases.html](https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html)

The following are the steps involved in the deploy stage configuration that will meet the requirements:

Use AWS CodeBuild to add sample event payloads for testing to the Lambda functions.

Publish a new version of the functions, and include Amazon CloudWatch alarms.

Update the production alias to point to the new version.

Configure rollbacks to occur when an alarm is in the ALARM state.

This configuration will help to reduce the customer impact of an unsuccessful deployment by deploying the new version of the functions to a staging environment first. This will allow the DevOps engineer to test the new version of the functions before deploying it to production.

The configuration will also help to monitor for issues by including Amazon CloudWatch alarms. These alarms will alert the DevOps engineer if there are any problems with the new version of the functions.

NEW QUESTION: 42

Which of the following configurations meets the requirements?

A. AWS Organizations defines the organizational structure. AWS CloudFormation defines the organizational structure.

B. AWS Organizations defines the organizational structure. AWS CloudFormation defines the organizational structure.

C. Active Directory defines the organizational structure. AWS CloudFormation defines the organizational structure.

D. AWS Organizations defines the organizational structure. AWS CloudFormation defines the organizational structure.

Which of the following configurations meets the requirements?

A. AWS Organizations defines the organizational structure. AWS CloudFormation StackSet defines the organizational structure. AWS IAM ID defines the organizational structure.

B. AWS Organizations defines the organizational structure. AWS CloudFormation StackSet defines the organizational structure. AWS IAM ID defines the organizational structure.

C. AWS Organizations defines the organizational structure. AWS Resource Access Manager(AWS RAM) defines the organizational structure. AWS IAM Identity Center defines the organizational structure.

D. AWS Organizations defines the organizational structure. AWS CloudFormation StackSet defines the organizational structure. AWS IAM ID defines the organizational structure.

Answer: D (LEAVE A REPLY)

This is a classic AWS Organizations governance design:

- * Restrict Regions and allowed services across all accounts
- * The correct Organizations mechanism for account-wide guardrails is a Service Control Policy (SCP) .
- * SCPs set the maximum permissions that accounts/OUs can use, making them the right tool to enforce "only these Regions" and "only these services" consistently across the org.
- * Authentication from Active Directory + consistent job-function roles in every account
- * With AD as the identity source, the standard approach is federation to AWS using an IAM identity provider (SAML) and role assumption based on job function.
- * AWS CloudFormation StackSets is the operationally efficient way to deploy identical IAM roles /policies into multiple accounts/OUs so job-function permissions are consistent everywhere.
- * The roles' trust policies can be configured to allow federated identities (from the AD-backed IdP) to assume the correct job role in each account.

Why the other options don't fit:

- * A "OU with group policies" isn't an AWS Organizations control mechanism for restricting Regions /services. The AWS-native control for this is SCPs , not "group policies."
- * B Permissions boundaries are useful to limit what an IAM principal/role can do within an account , but they are not the org-wide enforcement mechanism for "all accounts must stay in these Regions / use only these services." That's what SCPs do.
- * C AWS RAM is for sharing resources (like subnets, Transit Gateway, etc.), not for "sharing roles" across accounts as a governance baseline. Also, the question emphasizes identical permissions in each account, which is best achieved by provisioning roles in each account (StackSets), not attempting to "share" roles.

NEW QUESTION: 43

DevOps engineers are migrating an application from on-premises to AWS. The application consists of a MySQL database and a Lambda function. The database is currently hosted on Amazon Aurora. The Lambda function is currently hosted on Amazon EventBridge. The DevOps engineer wants to improve the performance of the application by using Amazon RDS Proxy. The engineer is considering the following options:

- A. Amazon RDS Proxy with Aurora. This option would allow the Lambda function to connect to the database through the proxy, which would reduce the number of connections to the database and improve performance.
- B. Lambda with EventBridge. This option would allow the Lambda function to connect to the database through the proxy, which would reduce the number of connections to the database and improve performance.
- C. Lambda with RDS Proxy. This option would allow the Lambda function to connect to the database through the proxy, which would reduce the number of connections to the database and improve performance.
- D. Lambda with Aurora. This option would allow the Lambda function to connect to the database through the proxy, which would reduce the number of connections to the database and improve performance.
- E. Lambda with RDS Proxy and Aurora. This option would allow the Lambda function to connect to the database through the proxy, which would reduce the number of connections to the database and improve performance.
- F. Lambda with Aurora and EventBridge. This option would allow the Lambda function to connect to the database through the proxy, which would reduce the number of connections to the database and improve performance.

Answer: A,C,E (LEAVE A REPLY)

Short 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. 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:

Implementing database connection pooling inside the Lambda code is unnecessary and redundant when using Amazon RDS Proxy, which already provides connection pooling as a service.

Implementing the database connection opening and closing inside the Lambda event handler code is inefficient and costly, as it can increase latency and consume more resources for each invocation of the Lambda function.

Connecting to the Aurora cluster endpoint from the Lambda function is not optimal for read-only queries, as it can direct traffic to either the primary instance or one of the Aurora Replicas in the DB cluster. This can result in inconsistent performance and potential conflicts with write operations on the primary instance.

NEW QUESTION: 44

DevOps is planning to deploy an application to Amazon EC2 instances using Amazon Elastic Beanstalk. The application uses a database connection pool. The DevOps engineer wants to ensure that the application can handle a high volume of traffic. Which of the following is the most appropriate configuration for the application?

A. ALB with 10 EC2 instances in a single Availability Zone, using Amazon Route 53 for DNS.

B. 10 EC2 instances in a single Availability Zone, using Amazon Route 53 for DNS.

C. 10 EC2 instances in a single Availability Zone, using Amazon CodeDeploy for deployment.

D. 10 EC2 instances in a single Availability Zone, using Amazon Elastic Beanstalk with an .ebextensions file to configure the application.

Answer: C (LEAVE A REPLY)

https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueInstanceTerminationOption.html The original revision termination settings are configured to wait 1 hour after traffic has been rerouted before terminating the blue task set.

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/deployment-type-bluegreen.html>

NEW QUESTION: 45

Which AWS CodePipeline action provides the most control over the production deployment of an application?
 A. AWS CodeConnections with Git
 B. Manual Approval
 C. AWS Lambda
 D. AWS CodeConnections with Git

- A. AWS CodeConnections with Git requires a new repository, a new service, and custom code to write, run, secure, and maintain.
- B. Manual Approval requires no new repositories, no new services, and no custom code—just pipeline configuration.
- C. AWS Lambda requires custom Lambda logic to inspect templates and decide whether to proceed—more code to write, run, secure, and maintain.
- D. AWS CodeConnections with Git requires a new repository, a new service, and custom code to write, run, secure, and maintain.

Answer: (SHOW ANSWER)

The requirement is simply: review changes before production deployment, with the least operational overhead.

* B is the lightest change: adding a Manual approval (review) action in CodePipeline creates a controlled gate before the deploy stage. It requires no new repositories, no new services, and no custom code—just pipeline configuration.

Why not the others:

- * A introduces additional moving parts (Git repo integration, PR workflow management, and CloudFormation Git sync). That's useful, but it's more operational overhead than necessary to satisfy "review before deploy."
- * C requires custom Lambda logic to inspect templates and decide whether to proceed—more code to write, run, secure, and maintain.
- * D adds both Git integration and a manual approval step—again more overhead than just adding the approval gate.

So B best meets the requirement with the least operational effort: a simple manual approval stage in the pipeline before production deployment.

NEW QUESTION: 46

Which IAM policy will allow a user to create a new Amazon EC2 instance?
 A. Allow all actions on all resources
 B. Allow Amazon RDS actions
 C. Allow IAM actions
 D. Allow AWS IAM actions

- A. Allow all actions on all resources allows all actions on all resources.
 - B. Allow Amazon RDS actions allows Amazon RDS actions.
 - C. Allow IAM actions allows IAM actions.
 - D. Allow AWS IAM actions allows AWS IAM actions.
- Which IAM policy will allow a user to create a new Amazon EC2 instance?
 A. Allow all actions on all resources
 B. Allow IAM actions
 C. Allow ResourceTag PrincipalTag StringEquals
 D. Allow IAM actions

AWS X-Ray provides distributed tracing for microservice-based applications. Deploying the X-Ray daemon as a DaemonSet in the EKS cluster and instrumenting the application with the X-Ray SDK enables end-to-end tracing across microservices, helping identify performance bottlenecks. This method is documented in "Using AWS X-Ray with Amazon EKS" (AWS Observability Guide).

NEW QUESTION: 48

A company is migrating a MySQL database to Amazon Aurora. The company requires a low RPO and RTO for the data. Which two deployment strategies should the company use to meet these requirements? (2 correct answers.)

- A. Create an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data store. Create an Amazon Aurora Global Database in two AWS Regions as the data store.
- B. Create an Amazon Aurora Global Database in two AWS Regions as the data store. Create an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data store.
- C. Create an Amazon Aurora Global Database in two AWS Regions as the data store. Create a Network Load Balancer in each Region and use it to route traffic to the Amazon Aurora Global Database.
- D. Create an Amazon Aurora Global Database in two AWS Regions as the data store. Create an Application Load Balancer in each Region and use it to route traffic to the Amazon Aurora Global Database.
- E. Create an Amazon Aurora Global Database in two AWS Regions as the data store. Create an Application Load Balancer (ALB) in each Region and use it to route traffic to the Amazon Aurora Global Database.

Answer: (SHOW ANSWER)

To meet the requirements of failover and disaster recovery, the company should use the following deployment strategies:

Create an Amazon Aurora global database in two AWS Regions as the data store. In the event of a failure, promote the secondary Region to the primary for the application. Update the application to use the Aurora cluster endpoint in the secondary Region. This strategy can provide a low RPO and RTO for the data, as Aurora global database replicates data with minimal latency across Regions and allows fast and easy failover¹². The company can use the Amazon Aurora cluster endpoint to connect to the current primary DB cluster without needing to change any application code¹.

Set up the application in two AWS Regions. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Regions. Add both ALBs to a single endpoint group. Use health checks and Auto Scaling groups in each Region. This strategy can provide high availability and performance for the application, as AWS Global Accelerator uses the AWS global network to route traffic to the closest healthy endpoint³. The company can also use static IP addresses that are assigned by Global Accelerator as a fixed entry point for their application¹. By using health checks and Auto Scaling groups, the company can ensure that their application can scale up or down based on demand and handle any instance failures⁴.

The other options are incorrect because:

Creating an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data store would not provide a fast failover or disaster recovery solution, as the company would need to manually restore data from backups or snapshots in another Region in case of a failure.

Creating an Amazon Aurora cluster in multiple AWS Regions as the data store and using a Network Load Balancer to balance the database traffic in different Regions would not work, as Network Load Balancers do not support cross-Region routing. Moreover, this strategy would not provide a consistent view of the data across Regions, as Aurora clusters do not replicate data automatically between Regions unless they are part of a global database.

Setting up the application in two AWS Regions and using Amazon Route 53 failover routing that points to Application Load Balancers in both Regions would not provide a low RTO, as Route 53 failover routing relies on DNS resolution, which can take time to propagate changes across different DNS servers and clients.

Moreover, this strategy would not provide deterministic routing, as Route 53 failover routing depends on DNS caching behavior, which can vary depending on different factors.

NEW QUESTION: 49

Auto Scaling □□□ Amazon EC2 □□□□ □□□□ □□□□ □□□□□□□□ □□□□□ □□ □□□ □□□□□. □□□□□ AWS CloudFormation□□ □□ □ □□ □□□□□. DevOps □□□□□ □□□□□ □□□ □ □□ □□ □□□ □□□ □□□ □□ □□□ □□□□□ □□□□□ □□□□□ □□□□□ □□□□□ □□□□□. □□ □□□ □□ □□ □□□□ □□ □□□□ AWS □□□ □□ □□□ □□ □□ □□□□□ □□□.

□□ □□□□ □□ □□□□□?

- A. CloudFormaiion □□□□□ AWS Config □□□□ □□□□□. □□□ InputParameters □□□ □□ □□ □□□□ □□□□□ Scope □□ □ EC2 Auto Scaling □□□□□ □□□□□. □□□□□ AWS Systems Manager □□□ □□□ □□□ □□□□□ □□□□□ □□□□□.
- B. CloudFormation □□□□□ EC2 □□ □□□ □□□□□ □□□□□. □□ □□□□ □□ □□ □□□□ □□□□□. □□□□□ □□□ □□□□□ cfn-mit □□□□□ □□□□ □□ □□□□□ □□□□□ cfn-hup □□□□□ □□□□□.
- C. CloudFormation □□□□□ EC2 □□ □□□ □□□□□ □□□□□. □□ □□□□ □□ □□ □□□□ □□□□□. □□□□□ AWS Systems Manager □□□ □□□ □□□ □□□□ □□□□ □□ □□□□□ □□□□□.
- D. CloudFormation □□□□□ CloudFormation imt □□□□□□□ □□□□□□. □□ □□ □□□□ □□□□□□ □□□□□. □□□□□ □□□ □□□□□ cfn-init □□□□□ □□□□ □□ □□□□□ □□□□□ cfn-hup □□□□□ □□□□□.

Answer: D (LEAVE A REPLY)

Use the AWS::CloudFormation::Init type to include metadata on an Amazon EC2 instance for the cfn-init helper script. If your template calls the cfn-init script, the script looks for resource metadata rooted in the AWS::CloudFormation::Init metadata key. Reference:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-init.html>

NEW QUESTION: 50

□ □□□ □ □□□□□□□□ □□ □□□ CI/CD □□□□□□□□ AWS CodeBuild□ □□□□ □□□□□. □ □□□ □□ □□ □□□□□□ □□ □□□□□ □□□□□ □□□. □□ □ □□□□□□ □□ □□□□ □□□□□□□ □□□ □ □□□ □□□. □□□□□□ □□ □□ □ □ □□ □□ Amazon S3 □□□□ □□□□□□ □□□. □□□□ □□□□□ □□□□ □□□ Git □□□□□ □□□□□ □□□. □ □□□ □□ □□□ □□□□□□ □□□. □□□ □□ □□□ □□ □□□□□□ □□□□ □□□□□ □□□□□?

- A. buildspec.yml □□□ □□□□□ □□ □□ □□ □□□ □□ □□ □□□□□ □□□□□. □□ □□ □□□ □□□□ □□ □□□□□ □□ □□□ □□□□□.
- B. □ □□□□ □□□□□ □□□ CodeBuild □□□□□ □□□□□. Docker □□ □□ □□ □□□□ □□ □□□□□ □□ □□□ □□ □□□□□ □□ ID□ □□□□ □□ □□ □□□ □□□ □□ □□□□□ S3 □□□ □□□□□ □□□□□. AWS Step Functions□ □□□□ □□□□□ □□□ □□□□□.
- C. CodeBuild□ □□□□ □□ □□□ □□□□ □□ ZIP □□□□□□ □□□□□. □□ □□□□□ S3 □□□ □□□□□□□. AWS CodePipeline□□ □□□ □□ □□□ □□□□ □□□ □□□□□ S3 □□□ □□ □□□□□ □□□□□□□. □□ □□□ □□□□□ □□ □□□ □□□□□ □□□□□ □□□ □□ □□□ □□□□□.
- D. □□ □□□ □□ □□□□□ □□□□□□□ buildspec.yml □□□ □□□□□□. □□□□ □□ □□ □□□ □□□□ □□ □□□□□ □□ □□ □□□□□□□. CodeBuild □□ □ □□□□□ AWS Lambda □□□ □□□□□□. Lambda □□□ □□□□□□□ □□□□□□ □□ □□ □ □□□□ □□ ID□ □□□□ □□ □□ □□□ □□□□□ S3 □□□ □□□□□□□.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 51

A company is developing a new application that will be hosted on Amazon Elastic Kubernetes Service (Amazon EKS). The application is stateless and must be highly available. The application must be able to handle traffic from multiple regions. The application must be able to handle traffic from multiple Availability Zones within each region. The application must be able to handle traffic from multiple Availability Zones across multiple regions. The application must be able to handle traffic from multiple Availability Zones across multiple regions. The application must be able to handle traffic from multiple Availability Zones across multiple regions. The application must be able to handle traffic from multiple Availability Zones across multiple regions.

- A. Deploy the application to a single Availability Zone in a single region. Use Amazon Route 53 to route traffic to the application. Use Amazon Route 53 to route traffic to the application. Use Amazon Route 53 to route traffic to the application.
- B. Deploy the application to a single Availability Zone in a single region. Use Amazon Route 53 to route traffic to the application. Use Amazon Route 53 to route traffic to the application. Use Amazon Route 53 to route traffic to the application.
- C. Deploy the application to multiple Availability Zones in multiple regions. Use Amazon Route 53 to route traffic to the application. Use Amazon Route 53 to route traffic to the application. Use Amazon Route 53 to route traffic to the application.
- D. Deploy the application to multiple Availability Zones in multiple regions. Use Amazon Route 53 to route traffic to the application. Use Amazon Route 53 to route traffic to the application. Use Amazon Route 53 to route traffic to the application.

Answer: C ([LEAVE A REPLY](#))

The requirement is to deploy a stateless application with multi-Region fault tolerance, ensuring high availability even if an entire AWS Region becomes unavailable. For this design, traffic must be actively served from both Regions, not only during a failure event. Option C correctly implements an active-active, multi-Region architecture. By deploying the application to both EKS clusters, each Region is capable of serving traffic independently. Using Amazon Route 53 weighted routing, traffic is distributed across both Application Load Balancers, allowing both Regions to handle requests simultaneously. If one Region becomes unhealthy, Route 53 health checks can stop routing traffic to that Region, maintaining availability. Implementing Kubernetes readiness and liveness probes ensures that traffic is only sent to healthy pods within each cluster. This provides fault tolerance at both the container level (pod health) and the Regional level (Route 53 routing). Option A uses a failover routing policy, which results in an active-passive design. While fault tolerant, it does not utilize both Regions simultaneously and provides slower recovery during Region failure. Options B and D deploy the application only in the primary Region, which does not meet multi-Region fault tolerance requirements. Therefore, Option C delivers the most resilient, highly available, and AWS-recommended architecture for a stateless, multi-Region EKS application.

NEW QUESTION: 52

A company is developing a new application that will be hosted on Amazon CloudFront. The application is stateless and must be highly available. The application must be able to handle traffic from multiple regions. The application must be able to handle traffic from multiple Availability Zones within each region. The application must be able to handle traffic from multiple Availability Zones across multiple regions. The application must be able to handle traffic from multiple Availability Zones across multiple regions. The application must be able to handle traffic from multiple Availability Zones across multiple regions. The application must be able to handle traffic from multiple Availability Zones across multiple regions.

C. You can use CloudFront to restrict access to a specific IP address range. You can use an ACL to restrict access to a specific IP address range. You can use an ACL to restrict access to a specific IP address range. You can use an ACL to restrict access to a specific IP address range.

D. You can use IP address sets to restrict access to a specific IP address range. You can use an ACL to restrict access to a specific IP address range. You can use an ACL to restrict access to a specific IP address range. You can use an ACL to restrict access to a specific IP address range.

Answer: D (LEAVE A REPLY)

To restrict access to CloudFront to a specific IP address range:

* Create an AWS WAF IP address set with the corporate office IPs.

* Modify the existing WebACL's default action to Block (deny all except explicitly allowed).

* Add a high-priority rule that allows traffic from the IP address set (the corporate IPs). This way, only requests from the corporate IPs are allowed; all others are blocked. Regex pattern sets are not necessary for IP-based restrictions and add complexity. Setting default action to Allow with exceptions is less secure and more complex to manage.

References:

AWS WAF IP Set Examples

Restricting Access by IP Address

NEW QUESTION: 53

You need to integrate AWS CloudTrail with Amazon EC2 instances in your AWS environment. You want to ensure that CloudTrail logs are sent to a Slack channel in your DevOps team's Slack workspace. How can you accomplish this?

Options:

A. Use the AWS Trusted Advisor to check the configuration of AWS Config and AWS Lambda. Use the Amazon Simple Notification Service (Amazon SNS) to send notifications to Slack.

B. Use Amazon EventBridge to send events to AWS Lambda, which can then send notifications to Slack via the Amazon SNS.

C. Use EC2 instance profiles to send events to Slack via the AWS Lambda and Amazon CloudWatch.

D. Use AWS CloudTrail to send events to AWS Lambda, which can then send notifications to Slack via the Amazon SNS.

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 54

You need to integrate AWS CloudTrail with Amazon Linux instances in your AWS environment. You want to ensure that CloudTrail logs are sent to a Slack channel in your DevOps team's Slack workspace. How can you accomplish this?

Options:

A. Use the AWS Trusted Advisor to check the configuration of AWS Config and AWS Lambda. Use the Amazon Simple Notification Service (Amazon SNS) to send notifications to Slack.

B. Use Amazon EventBridge to send events to AWS Lambda, which can then send notifications to Slack via the Amazon SNS.

C. Use EC2 instance profiles to send events to Slack via the AWS Lambda and Amazon CloudWatch.

D. Use AWS CloudTrail to send events to AWS Lambda, which can then send notifications to Slack via the Amazon SNS.

Which of the following is the most efficient and resilient design for a multi-region application that requires high availability and low latency? (3 correct answers)

Which of the following is the most efficient and resilient design for a multi-region application that requires high availability and low latency? (3 correct answers)

A. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2.

B. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

C. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

D. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

E. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

F. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

Answer: A,C,D (LEAVE A REPLY)

* (A) Use ECR replication to keep images synchronized between Regions, minimizing CI/CD pipeline changes.

* (C) DynamoDB global tables allow multi-Region replication and provide local read/write access, so tasks should interact with the DynamoDB replica in their Region.

* (D) Use S3 cross-Region replication or separate buckets with replication; tasks access the bucket in the same Region for latency and data sovereignty.

* (B) Using DynamoDB global table but pointing tasks to only one Region reduces resilience.

* (E) S3 Multi-Region Access Points are newer but add complexity and are not required for minimal changes.

* (F) Managing CI/CD pipeline for multiple Regions and failover routing adds complexity beyond minimal changes.

References:

Amazon ECR Cross-Region Replication

DynamoDB Global Tables

Amazon S3 Cross-Region Replication

NEW QUESTION: 59

Which of the following is the most efficient and resilient design for a multi-region application that requires high availability and low latency? (3 correct answers)

Which of the following is the most efficient and resilient design for a multi-region application that requires high availability and low latency? (3 correct answers)

A. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2.

B. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

C. AmazonEC2ReadOnlyAccess IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. AmazonEC2ReadOnlyAccess IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

D. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

E. Create an Amazon EC2 Read Only Access IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. Create an Amazon EC2 Read Only Access IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

F. AmazonEC2ReadOnlyAccess IAM role in eu-west-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in eu-west-2. AmazonEC2ReadOnlyAccess IAM role in ap-southeast-2 and use it to access Amazon S3, Amazon DynamoDB, and Amazon Route 53 in ap-southeast-2.

Answer: (SHOW ANSWER)

(A) When Docker communicates with an Amazon Elastic Container Registry (ECR) repository, it requires authentication. You can authenticate your Docker client to the Amazon ECR registry with the help of the AWS CLI (Command Line Interface). Specifically, you can use the "aws ecr get-login-password" command to get an authorization token and then use Docker's "docker login" command with that token to authenticate to the registry. You would need to perform these steps in your buildspec.yml file before attempting to push or pull images from/to the ECR repository.

DOP-C02-KR □□ □□□ □□□□□ □□ DumpTop □□ □□□□ □□□ DOP-C02-KR □□! DumpTop □ □□ **DOP-C02-KR** □□ □□□ □□□□□□, DumpTop DOP-C02-KR □□ □□□ □□□□□□□□ □□□ □□□□□□□□. □□□□ □□□ □□□□ □□ □□□□□□□□□□□□. <https://www.dumptop.com/Amazon/DOP-C02-KR-dump.html> (439 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)

NEW QUESTION: 62

□□□□ □□ AWS □□□ □□□□ VPC□ □□□□□□□ □□□□□□□. □ □□ □□□□□ □□□ □□□ □□□□ □□□□. □□□ □□ □□ AWS WAF□ □□ □□ □□□ □□□□□□ □□□ □□□□□ □□□. □□□ □□□□□□□ □□ □□□□ □□□ □□□□ □□ □□ □□ □□□ □□□□□ □□ □□□ □ □□ □ □□□ □□□□□ □□□ □□□□.

□□ □□ □□□□□□ □□ □ □□□ □□ □□□ □□ IP □□□ □□ □□□□□□ □ □ □□□ □□□□□. □□ □□ DevOps □□ □□□□ □□ □□□ □□□□□. □□ □□□ □□ IP □□ □ □□□□ □□□□□□□ □□□□□ □□ □□ □□ □□ □□□□□ □□ □□ □□□ □□□□ □ □□□ □□ □□□□□ □□□□ □□□ □□□ □□□□. DevOps □□□□□ □□□□ VPC□ □□ VPC □□ □□□ □□□□□.

□□□ □□ □□□ □□ □□ □□□□□ □□□□ □□ DevOps □□□□□ □□□□ □□ □□ □□□ □□□□□?

A. Amazon CloudWatch Logs□□ □□ □□□ □□□□□. □□□ □□□□ □□□□ □□□□ □□ □□□□ □□□□ VPC □□ □□□ □□□□□. □□ □□□ IP □□□ □□ Amazon CloudWatch □□ □□□ □□□□□. □□ □□□ □□□□ □□□□ CloudWatch □□ □□□□□. □□□ 5□□□□ □□□□ □□□ □□□ □□□□ 1□ □□□□□□. Amazon Simple Notification Service(Amazon SNS) □□ □□ □□□□ □□ □□□ □□ □□ □□□□.

B. □□ □□□ Amazon S3 □□□ □□□□□. □□□ □□□□ □□□□ □□□□ S3 □□□□ □□□□□ VPC □□ □□□ □□□□ □. □□ □□□ □□ Amazon OpenSearch Service □□□□ □ □□□□ □□□□□. S3 □□□□□ □□□ □□□□□, □□ □□□ □□□ □, OpenSearch Service □□□□□ □□□ □□□□□ AWS Lambda □□□ □□□□□. 5□□□ □□□□□□ Lambda □□□ □□□□□. □□ □□□ IP □□□□ □□□□ □□□□ Amazon Simple Notification Service(Amazon SNS) □□□ □□ □□ □□ □□□ □□□□ OpenSearch Service□□ □□ □ □□□ □□□□□.

C. Amazon CloudWatch Logs□□ □□ □□□ □□□□□. □□ □□□ □□□ Amazon S3 □□□ □□□□□. □□ □□□□ □□□□ □□□□ □□ □□□□ □□□□ VPC □□ □□□ □□□□□. AWS Lambda□□ Amazon Athena CloudWatch □□□□ □□□□□. □ □□□ □□ □□□ □□□□□. □□ □□□ IP □□□□ □□□ □□ □□□□ □□□□□ □□□□ □□□ □□□ □□□□□ Athena□ □□□□□. □ □□□ S3 □□□ □□□ □ Amazon Simple Notification Service(Amazon SNS) □□□ □□ □□ □□ □□□□ □□□□ S3 □□□ □□□ □□□□□.

D. □□ □□□ Amazon S3 □□□ □□□□□. □□ □□□□ □□□□ □□□□ S3 □□□□ □□□□□ VPC □□ □□□ □□□□□. □□ □□□ IP □□□ □□ S3 □□□ □□ □□ □□□ □□□□□ Amazon Athena□ □□□□□. Athena□□ □□□□ □□□□ □□ □□□□ □ □□ □□□□□ □□□□ □□□□□ Amazon QuickSight□ □□□□□. □□□□ □□□□ □□ □□□ □□ 1□ □□ □□. □□ □□□□ □□□□ □□□ □ □□ □□ □□ □□□□ □□□□ □□□□ □□□□□.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 63

DevOps □□□□□ Auto Scaling □□□ Application Load Balancer □□□ □□□□ □□ Amazon EC2 □□□□□ □□□ □□□ □□□ □ □□□ □□ □□□□□□. EC2 □□□□□ □□ □□ HTTP □□ □□□ □□□□□. □□ □□ □□□□□ EC2 □□□□□□ □□□ □□□□□ □□□□ □□ □□□ □□□□□□. □□□□ □□□ □□□ □□ □□□□ □□□ □□ □□□□. □□□□□□ □□□□ □□□□□□ □□□□ □□□□ □□□. □□□□ □□ □ □□□□□ □□□□ □ □□□ □ □□□□ □□□. □□ □□ □□□ □□□ □□ □□□ □□□□□? (2□□ □□□□□.)

- A. □□ □□□□ □□ □□□ □□□ □□ □□□□□ □□□□□ Auto Scaling □□□ □□□□□.
- B. □□ □□ □□ □□ HealthCheckIntervalSeconds □□□□□ □□□□ □□ □□ □□□ □□□□.
- C. □□ □□ □□ □□□ HTTP□□ TCP□ □□□□ □□□□□□□□ □□ □□ □□ □□□ □□□ □ □□□ □□□□□.
- D. □□ Auto Scaling □□□ □□ Amazon CloudWatch □□□□ □□□ □□ □□□ □□□ □□ □□□ □□□□□□. □□□ □□□□ □□ □ □□□ □□□□□. □□□□ □□ □ □□□ □□□□ Amazon SNS □□□ □□□ □□□□□.
- E. Amazon CloudWatch □□□□□ □□□□□ Auto Scaling □□□ □□ EC2 □□□□□ □□□ □□□□ □□□□□. □□□ □□□□ □ □ □□□ □□□□□ Amazon SNS □□□ □□□□ □□□ □□□□.

Answer: A,E (LEAVE A REPLY)

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/metrics-collected-by-CloudWatch-agent.html>

NEW QUESTION: 64

DevOps □□□□□ AWS Control Tower□ □□□□ □□ □ □□□□, □□ □ □□ □□□ □□□□ □□ □□ AWS □□□ □□□□□. □□ □□ DevOps □□□□□□ □□□ □□□□□ AWS Control Tower □□□□□ □□□□ □ □□ □□□□□. □ □□□□□□ AWS Organizations□□ □□ □□□ OUs □□ □□ □□ □□□□□ □□□. □□□□ □□ □□ □□□ □□ □□□□ □□□□□. □□□□ □□□ □□ □□□ □□□□ □□□ □ □□□ □□□. □□□□ □□□ □□□(OU) □□□ □□□□ □□□ □□□□. □□□□ □□□□ □□ □□□ □□□ □□□□□ □□□□□ □□□. □□ □□□□ □□□ □□□ □□ □□□ □□ □□□□□ □□ □□□□.

- A. □□□ □ □□□□□ □□ AWS CloudFormation □□□□ □□□□□. □□□□□ AWS CodeConnections□ □□□□□ Git □□□□□□ □□□□□. □□□ □ OUs □□ □□□□□ AWS::ControlTower::EnableControl □□ □□□□ □□□□□. Git □□□□□□□ □□□□ □ □□□ □□□□□ AWS CodeBuild □□□□□ □□□□□.
- B. □□□ □ □□□□□ □□ □□ AWS CloudFormation □□□□ □□□□□. □□□□□ AWS CodeConnections□ □□□□□ Git □□□ □□□ □□□□□. □□ □ □ □□□ □□ □□□□□ AWS::ControlTower::EnableControl □□ □□□□ □□□□□. □□ □ □□□□ AWS CodePipeline□□ □□□□□□□ □□□□□. □□ □□ □□□ □□□ □ □□□□□□ □□□□ □□□□ □□□□ □□□□ □□□□ □□□□ □ □□□.
- C. □□□ □□□□□ □□ □□ AWS CloudFormation □□□□ □□□□□. □□□□□ AWS CodeConnections□ □□□□□ Git □□□□□ □ □□□□□. □□□ □ OUs □□ □□□□□ AWS::ControlTower::EnableControl □□ □□□□ □□□□□. □□ □ □□□□ AWS CodePipeline□□ □□□□□□□ □□□□□. □□ □□□ Git □□□□□□□ □□□ □□□ □ □□□□□□ □□□□□ Amazon EventBridge □□□ □□□□□.

Each project requires a separate code repository and a separate testing environment. Amazon S3 buckets can be used for code storage, but they do not provide the same level of revision control as AWS CodeCommit. Option B is incorrect because creating another S3 bucket for each project for testing code and using an AWS Lambda function to promote code changes between testing and production buckets will not provide the benefits of revision control, such as tracking changes, branching, merging, and collaborating. Option C is incorrect because using the main branch for production and test code with different deployment pipelines for each environment will not allow the developers to test their code changes before deploying them to production. Option D is incorrect because enabling versioning and branching on each S3 bucket will not work with Git-based tools and will not provide the same level of revision control as AWS CodeCommit. References:

Which of the following options best meets the requirements?

- A. Create an AWS CodeCommit repository for each project, use the main branch for production code, and create a testing branch for code deployed to testing.
- B. Create another S3 bucket for each project for testing code and use an AWS Lambda function to promote code changes between testing and production buckets.
- C. Use AWS CodeCommit for each project, use the main branch for production code, and use different deployment pipelines for each environment.
- D. Enable versioning and branching on each S3 bucket, use the main branch for production code, and use different deployment pipelines for each environment.

Answer: A (LEAVE A REPLY)

Creating an AWS CodeCommit repository for each project, using the main branch for production code, and creating a testing branch for code deployed to testing will meet the requirements. AWS CodeCommit is a managed revision control service that hosts Git repositories and works with all Git-based tools¹. By using feature branches to develop new features and pull requests to merge code to testing and main branches, the developers can avoid code conflicts and lost work, and also implement code reviews and approvals. Option B is incorrect because creating another S3 bucket for each project for testing code and using an AWS Lambda function to promote code changes between testing and production buckets will not provide the benefits of revision control, such as tracking changes, branching, merging, and collaborating. Option C is incorrect because using the main branch for production and test code with different deployment pipelines for each environment will not allow the developers to test their code changes before deploying them to production. Option D is incorrect because enabling versioning and branching on each S3 bucket will not work with Git-based tools and will not provide the same level of revision control as AWS CodeCommit. References:

* AWS CodeCommit

* Certified DevOps Engineer - Professional (DOP-C02) Study Guide (page 182)

NEW QUESTION: 67

A company wants to store code in Amazon S3 and use AWS CodePipeline to build and deploy the code to Amazon Elastic Container Service (Amazon ECS) using Amazon Elastic Container Registry (Amazon ECR). Which of the following options best meets the requirements?

Which of the following options best meets the requirements?

Which of the following options best meets the requirements?

Which of the following options best meets the requirements?

- A. Amazon ECR VPC, Amazon S3, and AWS CodePipeline. Use AWS CodePipeline to build and deploy the code to Amazon ECS using Amazon ECR.
- B. Amazon S3, Amazon ECR, and AWS CodePipeline. Use AWS CodePipeline to build and deploy the code to Amazon ECS using Amazon ECR.

- B. Use Amazon S3 to store code, use Amazon ECR to store Docker images, and use AWS CodePipeline to build and deploy the code to Amazon ECS using Amazon ECR.

The other options are not optimal because they either require more development overhead, do not provide near real time detection and remediation, or do not provide a centralized dashboard for tracking.

Option A is not optimal because CloudFormation drift detection is not a near real time solution. Drift detection has to be manually initiated on each stack or resource, or scheduled using a cron expression. Drift detection also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. CloudWatch Logs and dashboard can be used for tracking, but they do not provide a comprehensive view of the security posture of the AWS accounts and resources.

Option B is not optimal because CloudTrail logs analysis using Athena is not a near real time solution. Athena queries have to be manually run or scheduled using a cron expression. Athena also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. Step Functions can be used to orchestrate the query and remediation workflow, but it adds more complexity and cost. QuickSight dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.

Option D is not optimal because CloudTrail logs analysis using CloudWatch Logs is not a near real time solution. CloudWatch Logs filters have to be manually created or updated for each resource type and configuration change. CloudWatch Logs also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. EventBridge can be used to trigger the Lambda function, but it adds more complexity and cost. OpenSearch Service dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.

AWS Config conformance packs

Introducing AWS Config conformance packs

Managing conformance packs across all accounts in your organization

NEW QUESTION: 71

□□□□ AWS□ □ □□□□□□□ □□□□ □□□□. □□□□□□□□ Application Load Balancer □□ Amazon EC2 □□□□□□ □□ □□□. □□□□□□ □□ □□ □□□□ Auto Scaling □□□□ □□□□□□. □□□□□□□□ Oracle DB □□□□□ □ Amazon DynamoDB□ Amazon RDS□ □□□□□ □□□□□□. □□ □□□□□ □□□□□□□ □□□ □□□ □□□□□. □□ □□ □□ □□ □□□ □□ □ □□ □□ □□□□ □□□ □□□ □□□□□□?

A. AWS □□□□ □□□□□ □□ AWS Systems Manager securestring □□□□□□ □□□ □□ □□□□□□. Systems Manager SecureString □□□□□□ □□□□□□ □□ □□□□ □□□□□□.

B. EC2 1AM □□□□ EC2 □□□□□□ □□□□□□ AWS □□□□ □□□□□□□. AWS Secrets Manager□□ □□□□□□□ □□ □□□□ □□ □□□.

C. AWS □□□□ □□□□□ □□ AWS Systems Manager □□ □□□ □□□□□□ □□□ □□ □□□□□□. Systems Manager SecureString □□□□□□ □□□□□□ □□ □□□□ □□□□□□.

D. EC2 1AM □□□□ EC2 □□□□□□ □□□□□□ AWS □□□□ □□□□□□□. □□□□□□□ □□□□□□ □□ □□□□ □□ □□□□ □□□□ □□□□□□.

Answer: **B** ([LEAVE A REPLY](#))

AWS Secrets Manager is a secrets management service that helps you protect access to your applications, services, and IT resources. This service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Using Secrets Manager, you can secure and manage secrets used to access resources in the AWS Cloud, on third-party services, and on-premises. SSM parameter store and AWS Secret manager are both a secure option. However, Secrets manager is more flexible and has more options like password generation. Reference: <https://www.1strategy.com/blog/2019/02/28/aws-parameter-store-vs-aws-secrets-manager/>

NEW QUESTION: 72

Which of the following AWS services can be used to manage the lifecycle of an Amazon EC2 instance? (Select two.)

A. AWS Auto Scaling Group, Amazon EC2 Instance Profile, Application Load Balancer (ALB)

B. AWS CloudFormation Stack, Amazon EC2 Instance Profile, Application Load Balancer (ALB)

C. AWS CloudFormation Stack, Amazon EC2 Instance Profile, Amazon S3 Bucket

D. AWS CloudFormation Stack, Amazon EC2 Instance Profile, Amazon S3 Bucket

- A. ALB and Amazon EC2 Instance Profile. aws cloudformation update-stack-set AWS CLI
- B. Amazon EC2 Instance Profile and Application Load Balancer (ALB). aws cloudformation continue-update-rollback AWS CLI
- C. Amazon EC2 Instance Profile and Amazon S3 Bucket. aws cloudformation cancel-update-stack AWS CLI
- D. Auto Scaling Group and Amazon S3 Bucket. aws cloudformation rollback-stack AWS CLI

Answer: B (LEAVE A REPLY)

<https://repost.aws/knowledge-center/cloudformation-update-rollback-failed> If your stack is stuck in the UPDATE_ROLLBACK_FAILED state after a failed update, then the only actions that you can perform on the stack are the ContinueUpdateRollback or DeleteStack operations.

NEW QUESTION: 73

DevOps wants to use S3 to store application logs. Which of the following IAM roles can be used to grant permissions to the Amazon S3 bucket? (Select three.)

A. AmazonS3OutpostsRole

B. AmazonS3OutpostsRole

C. AmazonS3OutpostsRole

D. AmazonS3OutpostsRole

E. AmazonS3OutpostsRole

F. AmazonS3OutpostsRole

- A. AmazonS3OutpostsRole
- B. AmazonS3OutpostsRole
- C. AmazonS3OutpostsRole
- D. AmazonS3OutpostsRole
- E. AmazonS3OutpostsRole
- F. AmazonS3OutpostsRole

Answer: A,D,E (LEAVE A REPLY)

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: 74

Which of the following AWS services can be used to manage the lifecycle of an Amazon EC2 instance? (Select two.)

A. AWS CodeDeploy, Amazon EC2 Instance Profile, Application Load Balancer (ALB)

B. AWS CodeDeploy, Amazon EC2 Instance Profile, Amazon S3 Bucket

C. AWS CodeDeploy, Amazon EC2 Instance Profile, Amazon S3 Bucket

D. AWS CodeDeploy, Amazon EC2 Instance Profile, Amazon S3 Bucket

- A. Amazon EC2 Instance Profile and SSH. CodeDeploy

- B. EC2 Image Builder is the recommended AWS-native service for automating the creation, validation, and distribution of AMIs. It integrates directly with AWS Organizations, allowing AMIs to be automatically shared with all accounts. The service also supports scheduled builds (e.g., monthly), significantly reducing operational overhead. This design is referenced in AWS documentation: "Automate OS image creation and distribution securely across multiple accounts using EC2 Image Builder."
- C. Amazon EventBridge Scheduler is the recommended AWS-native service for automating the creation, validation, and distribution of AMIs. It integrates directly with AWS Organizations, allowing AMIs to be automatically shared with all accounts. The service also supports scheduled builds (e.g., monthly), significantly reducing operational overhead. This design is referenced in AWS documentation: "Automate OS image creation and distribution securely across multiple accounts using EC2 Image Builder."
- D. AWS Systems Manager Automation is the recommended AWS-native service for automating the creation, validation, and distribution of AMIs. It integrates directly with AWS Organizations, allowing AMIs to be automatically shared with all accounts. The service also supports scheduled builds (e.g., monthly), significantly reducing operational overhead. This design is referenced in AWS documentation: "Automate OS image creation and distribution securely across multiple accounts using EC2 Image Builder."

Answer: (SHOW ANSWER)

EC2 Image Builder is the recommended AWS-native service for automating the creation, validation, and distribution of AMIs. It integrates directly with AWS Organizations, allowing AMIs to be automatically shared with all accounts. The service also supports scheduled builds (e.g., monthly), significantly reducing operational overhead. This design is referenced in AWS documentation: "Automate OS image creation and distribution securely across multiple accounts using EC2 Image Builder."

DOP-C02-KR is a dump of the AWS DOP-C02-KR exam. DumpTop is the best place to buy DOP-C02-KR dumps. DumpTop DOP-C02-KR dumps are available for purchase. DumpTop DOP-C02-KR dumps are available for purchase. <https://www.dumptop.com/Amazon/DOP-C02-KR-dump.html> (439 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)

NEW QUESTION: 77

DevOps is a set of practices that combines software development and operations. The goal is to create a high-performance engineering culture by encouraging and enabling the closest collaboration between those who build and those who operate the software.

DevOps is a set of practices that combines software development and operations. The goal is to create a high-performance engineering culture by encouraging and enabling the closest collaboration between those who build and those who operate the software.

DevOps is a set of practices that combines software development and operations. The goal is to create a high-performance engineering culture by encouraging and enabling the closest collaboration between those who build and those who operate the software.

DevOps is a set of practices that combines software development and operations. The goal is to create a high-performance engineering culture by encouraging and enabling the closest collaboration between those who build and those who operate the software.

A. AWS CodePipeline is a fully managed service that automates your build and release pipeline. It integrates with AWS CodeConnections, Git, and Amazon ECS. CodeDeploy is a service that automates software deployments to a variety of compute services, including Amazon EC2, Amazon ECS, and AWS Lambda.

B. ECS is a managed service that runs your Docker containers on a fleet of EC2 instances. ALB is a managed service that routes traffic to your Amazon EC2 instances, Amazon ECS instances, or AWS Lambda functions. CodeDeploy is a service that automates software deployments to a variety of compute services, including Amazon EC2, Amazon ECS, and AWS Lambda.

C. AWS Lambda is a serverless compute service that runs your code in response to events. Amazon ECS RegisterTaskDefinition API is used to register task definitions for Amazon ECS. ECS is a managed service that runs your Docker containers on a fleet of EC2 instances.

D. AWS CodeBuild is a fully managed service that builds your source code. AWS CodeConnections is a service that connects your code repositories to your build services. Git is a distributed version control system. Amazon ECS UpdateService API is used to update the service configuration for Amazon ECS. ALB is a managed service that routes traffic to your Amazon EC2 instances, Amazon ECS instances, or AWS Lambda functions.

Answer: (SHOW ANSWER)

CodePipeline_S3_Bucket_Policy IAM Role DevOps_Role IAM Role PipelineAccount

CodePipeline_Service_Role IAM Role DevOps_Role IAM Role PipelineAccount

C. S3 Bucket Policy PipelineAccount IAM Role DevOps_Role IAM Role PipelineAccount IAM Role CodePipeline_Service_Role IAM Role DevOps_Role IAM Role

CodePipeline_Service_Role IAM Role DevOps_Role IAM Role

D. S3 Bucket Policy CodeDeployAccount IAM Role DevOps_Role IAM Role

CodeDeployAccount IAM Role CodePipeline_Service_Role IAM Role DevOps_Role IAM Role

Answer: A (LEAVE A REPLY)

For cross-account deployments, CodePipeline must assume a role in the target account. Configure S3 bucket policy to allow the target account (CodeDeployAccount) access, trust relationship in DevOps_Role to allow assumption by the pipeline's account (PipelineAccount), and update CodePipeline_Service_Role with sts:

AssumeRole permissions. This cross-account pipeline setup is documented in "Cross-Account Access for AWS CodePipeline."

NEW QUESTION: 80

A company is migrating its on-premises Windows and Linux applications to AWS and creating a pilot light DR environment in another AWS Region. The company wants to use Amazon EC2 instances in the DR Region to access the application data stored in the primary Region. Which solution meets these requirements?

A. Use Amazon S3 for application data storage. Configure EC2 instances in the DR Region to access the data via the AWS CLI.

B. Use Amazon Elastic Block Store (Amazon EBS) for application data storage. Configure EC2 instances in the DR Region to access the data via the AWS CLI.

C. Use Amazon S3 for application data storage. Configure EC2 instances in the DR Region to access the data via the AWS CLI.

D. Use Amazon Elastic Block Store (Amazon EBS) for application data storage. Configure EC2 instances in the DR Region to access the data via the AWS CLI.

E. Use Amazon Storage Gateway for application data storage. Configure EC2 instances in the DR Region to access the data via the AWS CLI.

F. Use Amazon FSx for NetApp ONTAP for application data storage. Configure EC2 instances in the DR Region to access the data via the AWS CLI.

Answer: D (LEAVE A REPLY)

To meet the requirements of migrating its on-premises Windows and Linux applications to AWS and creating a pilot light DR environment in another AWS Region, the company should use Amazon FSx for NetApp ONTAP for the application storage. Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx for ONTAP supports multiple protocols, including SMB for Windows and NFS for Linux, so the company can access the shared storage from both types of applications. FSx for ONTAP also supports NetApp SnapMirror replication, which enables the company to replicate the storage to the DR Region. NetApp SnapMirror replication is efficient, secure, and incremental, and it preserves the data deduplication and compression benefits of FSx for ONTAP. The company can use automation to launch and configure the EC2 instances in the DR Region and then use NetApp SnapMirror to restore the data from the primary Region.

The other options are not correct because they do not meet the requirements or follow best practices. Using Amazon S3 for the application storage is not a good option because S3 is an object storage service that does not support SMB or NFS protocols natively. The company would need to use additional services or software to mount S3 buckets as file systems, which would add complexity and cost. Using

Amazon EBS for the application storage is also not a good option because EBS is a block storage service that does not support SMB or NFS protocols natively. The company would need to set up and manage file servers on EC2 instances to provide shared access to the EBS volumes, which would add overhead and maintenance. Using a Volume Gateway in AWS Storage Gateway for the application storage is not a valid option because Volume Gateway does not support SMB protocol. Volume Gateway only supports iSCSI protocol, which means that only Linux applications can access the shared storage.

1: What is Amazon FSx for NetApp ONTAP? - FSx for ONTAP

2: Amazon FSx for NetApp ONTAP

3: Amazon FSx for NetApp ONTAP | NetApp

4: AWS Announces General Availability of Amazon FSx for NetApp ONTAP

Replicating Data with NetApp SnapMirror - FSx for ONTAP

What Is Amazon S3? - Amazon Simple Storage Service

What Is Amazon Elastic Block Store (Amazon EBS)? - Amazon Elastic Compute Cloud What Is AWS Storage Gateway? - AWS Storage Gateway

NEW QUESTION: 81

DevOps □□□□□ AWS CodePipeline□ □□□□ □□□ □□□□□□ □□□□ □□□□□□□ □□, □□, □□□□, □□□ □ □□ □□□. □□□ □□□ □□ □□ □□□ □□ □□ □□□ □□□□□. □□ □□ □□ □□□ □□□ □□□ □□□ □□□ □□ □□□ □□ □□□□□.

DevOps □□□□□ □□ □□□ □□□ □□□□□ □□ □ □□ □□□ □□ □□ □□□□□ □□□ □□□□ □□□?

A. CodePipeline □□□□□ □□ □□ □□□ □□□□□ Amazon CloudWatch Logs □□□ □□□□□. Amazon Simple Notification Service(Amazon SNS) □□□ □□ □□□□ □□□□□. SNS □□□ □□ □□ □□ URL□ □□□□ □□ □□□ □□□ □□□□□.

B. AWS CloudTrail □□□□ □□ □□□□ AWS Lambda □□□ □□□□□. CodePipeline □□□□□ □□ □□ □□ □□□□ □□□ □ □□□ □□ □□□ □□ □□□ URL□ □□□□.

C. CodePipeline □□□□□ □□ □□ □□□ □□□□□ Amazon EventBridge □□□ □□□□□. Amazon Simple Notification Service(Amazon SNS) □□□ □□□□ □□□□□. □□□ □□ □□□ □□ □□□ URL□ □□□ AWS Lambda □□□ □□□□□. SNS □□□ □□ □□□ □□□□□□.

D. □ □□ □□ □ □□ □□ URL□ □□□ □□□ □□□□ □□□□□ □□□ □□□□□. □ □□□□□□ □□□□□ □□□ □□ □ □ □□□ □□ □ □□□ □□□ □□□□□□□□.

Answer: (SHOW ANSWER)

<https://aws.amazon.com/premiumsupport/knowledge-center/sns-lambda-webhooks-chime-slack-teams/>

NEW QUESTION: 82

□ □□□ □□ □□□ AWS Systems Manager □□□□ □□□□ □□ □□ □□ □□□□□ □□□□□ □□□□ □□□□. □□ □□ □□ 60□□□□ □□□□ □□□□□ □□□.

DevOps □□□□□ □□□□□□□ Amazon ElastiCache(Redis OSS) □□□□□ □□□□ □ □□□ □ □□ □□ □□□□ □□□. □□ □□□□ □□ □□□ □□□□□□ □□□ □□ □□□ □□□□□?

A. AWS Secrets Manager□□ □□□□ □□□□□. □□□ □□□□□□. □□ □□□ 60□□ □□□□□. □□□□ □□□□□ □□ □□□ □□□□ □□□□□□□ □□ □□ □□□□□ □□□□□□.

B. □□□□ □□□□ □□ □□ □□□□□. □□ □□□ □□□□□□. □□ □□□ 60□□ □□□□□. □□□□ □□□□□ □□ □□□ □□□□ □□□□□□□ □□ □□ □□□□□ □□□□□□.

- C. Employees must access secrets and parameters through AWS Systems Manager Parameter Store. AWS Lambda functions use Amazon EventBridge to rotate secrets every 60 days. Lambda functions use Amazon EventBridge to rotate secrets every 60 days.
- D. AWS Secrets Manager supports automatic secret rotation and provides built-in rotation templates for supported services, including Amazon ElastiCache for Redis OSS. This gives the least operational overhead for rotating the secret every 60 days. Parameter Store can reference Secrets Manager secrets by using a reserved path, which allows the application to retrieve the secret through Parameter Store while the secret is actually stored and rotated in Secrets Manager.

Answer: D (LEAVE A REPLY)

Comprehensive and Detailed Explanation From Exact Extract of DevOps Engineer documents only:

The correct answer is D because AWS Secrets Manager supports automatic secret rotation and provides built-in rotation templates for supported services, including Amazon ElastiCache for Redis OSS. This gives the least operational overhead for rotating the secret every 60 days.

The requirement also states that employees and applications must access secrets and parameters through AWS Systems Manager Parameter Store. Parameter Store can reference Secrets Manager secrets by using a reserved path, which allows the application to retrieve the secret through Parameter Store while the secret is actually stored and rotated in Secrets Manager.

Why the other options are incorrect:

- A). This option correctly uses Secrets Manager rotation, but it does not specifically use the reserved Parameter Store reference path approach that satisfies the access requirement through Parameter Store as clearly as option D.
- B). Parameter Store does not provide native automatic rotation for secrets in the same way that Secrets Manager does.
- C). This option is possible, but it requires custom Lambda and EventBridge automation, which creates more operational overhead than using native Secrets Manager rotation with the provided template.

NEW QUESTION: 83

DevOps team is using AWS CodePipeline to build and deploy applications. A pipeline step is failing with a 503 HTTP error. The pipeline is using Amazon CloudWatch to monitor the pipeline. The pipeline is using AWS Lambda to run the pipeline. The pipeline is using Amazon API Gateway to expose the pipeline. The pipeline is using AWS Device Farm to test the pipeline. The pipeline is using Amazon S3 to store the pipeline artifacts. The pipeline is using Amazon ElastiCache to cache the pipeline artifacts. The pipeline is using Amazon IAM to manage the pipeline artifacts. The pipeline is using Amazon CloudFront to deliver the pipeline artifacts. The pipeline is using Amazon Rekognition to analyze the pipeline artifacts. The pipeline is using Amazon SageMaker to train the pipeline artifacts. The pipeline is using Amazon Forecast to predict the pipeline artifacts. The pipeline is using Amazon Personalize to recommend the pipeline artifacts. The pipeline is using Amazon Comprehend to analyze the pipeline artifacts. The pipeline is using Amazon Lex to chat with the pipeline artifacts. The pipeline is using Amazon Polly to synthesize the pipeline artifacts. The pipeline is using Amazon Textract to extract the pipeline artifacts. The pipeline is using Amazon Transcribe to transcribe the pipeline artifacts. The pipeline is using Amazon Translate to translate the pipeline artifacts. The pipeline is using Amazon Rekognition to analyze the pipeline artifacts. The pipeline is using Amazon SageMaker to train the pipeline artifacts. The pipeline is using Amazon Forecast to predict the pipeline artifacts. The pipeline is using Amazon Personalize to recommend the pipeline artifacts. The pipeline is using Amazon Comprehend to analyze the pipeline artifacts. The pipeline is using Amazon Lex to chat with the pipeline artifacts. The pipeline is using Amazon Polly to synthesize the pipeline artifacts. The pipeline is using Amazon Textract to extract the pipeline artifacts. The pipeline is using Amazon Transcribe to transcribe the pipeline artifacts. The pipeline is using Amazon Translate to translate the pipeline artifacts.

- A. AWS CodeDeploy is used to deploy the pipeline artifacts. CheckURL is used to check the pipeline artifacts. CodePipeline is used to build the pipeline artifacts.
- B. Amazon CloudWatch is used to monitor the pipeline artifacts. CheckURL is used to check the pipeline artifacts. CodePipeline is used to build the pipeline artifacts.
- C. URL is used to check the pipeline artifacts. CodePipeline is used to build the pipeline artifacts. AWS Lambda is used to run the pipeline artifacts. CheckURL is used to check the pipeline artifacts. Lambda is used to run the pipeline artifacts.
- D. URL is used to check the pipeline artifacts. CodePipeline is used to build the pipeline artifacts. Amazon API Gateway is used to expose the pipeline artifacts. CheckURL is used to check the pipeline artifacts. AWS Device Farm is used to test the pipeline artifacts. API Gateway is used to expose the pipeline artifacts.

Answer: (SHOW ANSWER)

NEW QUESTION: 84

How can you prevent developers from unintentionally attaching an Elastic IP address to an Amazon EC2 instance in production? Which approach is the best?

- A. Amazon Athena, AWS CloudTrail, AWS Lambda, IAM policies, and AWS Config rules.
- B. IAM policies, AWS Config rules, and Amazon CloudWatch alarms.
- C. IAM policies, AWS Config rules, and Amazon CloudWatch alarms.
- D. IAM policies, AWS Config rules, Amazon CloudWatch alarms, and Amazon CloudFront.

Answer: B (LEAVE A REPLY)

To prevent developers from unintentionally attaching an Elastic IP address to an Amazon EC2 instance in production, the best approach is to use IAM policies and AWS Config rules. By attaching an IAM policy that denies the associate-address permission to the developers' IAM group, you ensure that developers cannot perform this action. Additionally, creating a custom AWS Config rule to check for Elastic IP addresses associated with instances tagged as production provides ongoing monitoring. If the rule detects an Elastic IP address, it can trigger an alert to notify the security team. This method is proactive and enforces the necessary permissions while also providing a mechanism for detection and notification. References: from Amazon DevOps sources

NEW QUESTION: 85

You are using AWS Systems Manager to manage Amazon EC2 Linux instances. You want to ensure that all instances are configured with the same configuration. Which approach is the best?

- A. Use the AWS-ApplyChefRecipes Systems Manager document to apply Chef recipes to all instances.
- B. Use the AWS-InstallApplication Systems Manager document to install the Chef client on all instances.
- C. Use the AWS-RefreshAssociation Systems Manager document to refresh the configuration on all instances.
- D. Use the AWS-ConfigureInstance Systems Manager document to configure all instances.

Answer: A (LEAVE A REPLY)

Option A directly matches all requirements with the least extra infrastructure:
* State Manager is the Systems Manager capability for defining and maintaining a desired configuration by running associations on a schedule (for example, hourly). That satisfies the "run periodically" requirement in a managed way across a large fleet.
* The AWS-ApplyChefRecipes Systems Manager document is specifically intended to run Chef recipes

/cookbooks on managed instances without requiring you to run your own Chef server infrastructure . You can point it at a cookbook source (such as an artifact/repo location) and have Systems Manager handle execution on the instances.

Why the other options aren't correct:

* B is not the right mechanism for periodic Chef cookbook execution. Installing an application package and running a "repair action" via Run Command is not the purpose-built Chef cookbook runner, and it's more brittle/DIY than State Manager associations.

* C (AWS-RefreshAssociation) is used to refresh/reevaluate association metadata/targets; it does not itself execute Chef cookbooks as the compliance mechanism.

* D is for OS patching compliance (Patch Manager/patch policies), not for applying Chef cookbooks that detect and remediate configuration drift.

NEW QUESTION: 86

DevOps wants to ensure that AWS WAF is enabled on all Amazon Elastic Load Balancing (ALB) instances in the AWS CloudFormation stack. DevOps wants to ensure that AWS WAF is enabled on all ALB instances in the AWS CloudFormation stack. DevOps wants to ensure that AWS WAF is enabled on all ALB instances in the AWS CloudFormation stack. DevOps wants to ensure that AWS WAF is enabled on all ALB instances in the AWS CloudFormation stack.

Which of the following is the most operationally efficient and reliable approach to ensure continuous compliance?

A. AWS Config rule alb-waf-enabled in the AWS Config managed rules catalog. AWS Systems Manager Automation document that enables AWS WAF on ALB instances. Systems Manager Automation document that enables AWS WAF on ALB instances.

B. AWS Config rule alb-waf-enabled in the AWS Config managed rules catalog. AWS Config ConfigurationItemChangeNotification Lambda function that triggers an Amazon EventBridge event. Lambda function that triggers an AWS Config start-resource-evaluation API call.

C. CloudFormation stack detect-stack-drift API call. AWS Lambda function that triggers an Amazon EventBridge event. AWS Config rule AWS::WAFv2::WebACLAssociation that triggers a Lambda function. Lambda function that triggers an AWS Config start-resource-evaluation API call.

D. CloudFormation stack detect-stack-drift API call. AWS Lambda function that triggers an Amazon EventBridge event. AWS Config rule AWS::WAFv2::WebACLAssociation that triggers a Lambda function. Lambda function that triggers an AWS Config start-resource-evaluation API call.

Answer: A (LEAVE A REPLY)

AWS Config has a managed rule called alb-waf-enabled that checks whether AWS WAF is enabled on ALBs.

AWS Config supports automatic remediation actions that can be triggered when noncompliance is detected.

By creating a Systems Manager Automation document that adds AWS WAF to the ALB and associating it as the remediation action for the AWS Config rule, the system can automatically detect and remediate any removal of AWS WAF from ALBs without manual intervention.

This is the most operationally efficient and reliable approach to ensure continuous compliance.

Option B lacks automatic remediation. Options C and D rely on drift detection and Lambda, which add complexity and risk downtime during stack replacement.

Reference:

AWS Config Managed Rules: " The alb-waf-enabled rule checks for AWS WAF association with ALBs and supports automatic remediation using Systems Manager Automation. " (AWS Config Managed Rules) AWS Config Remediation: " AWS Config automatic remediation can invoke Systems Manager Automation documents to remediate noncompliance. " (AWS Config Remediation)

NEW QUESTION: 87

- 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.
- E. 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
- F. DynamoDB table records.

Answer: C,D,F (LEAVE A REPLY)

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: 90

- A DevOps engineer is responsible for ensuring that all Amazon EC2 instances are running on Amazon EC2 Dedicated Hosts. The engineer wants to create 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.
- A. AWS Systems Manager put-compliance-items API call to create a compliance item. The compliance item is associated with an EC2 instance profile. The compliance item is associated with an Amazon DynamoDB table. The compliance item is associated with an Amazon ID. The compliance item is associated with a list-compliance-summaries API call to Systems Manager.
 - B. EC2 instance profile associated with Java. The compliance item is associated with an EC2 Auto Scaling group. The compliance item is associated with an EC2 instance ID. The compliance item is associated with an Amazon SQS queue. The compliance item is associated with an Amazon DynamoDB table. The compliance item is associated with an Amazon SNS topic. The compliance item is associated with an Amazon ID.
 - C. AWS Config rule. The compliance item is associated with an Amazon EC2 instance profile. The compliance item is associated with an EC2 instance ID. The compliance item is associated with a "config-rule-change-triggered" event. The compliance item is associated with an AWS Lambda function. The compliance item is associated with an AWS Config rule. The compliance item is associated with an EC2 instance ID. The compliance item is associated with a NON_COMPLIANT result. The compliance item is associated with a Lambda function. The compliance item is associated with an AWS Config rule.
 - D. AWS CloudTrail. EC2 RunCommand API call to create a compliance item. The compliance item is associated with an EC2 instance profile. The compliance item is associated with an EC2 instance ID. The compliance item is associated with an Amazon SQS queue. The compliance item is associated with an Amazon RDS DB instance. The compliance item is associated with an Amazon RDS DB instance. The compliance item is associated with an Amazon CSV file. The compliance item is associated with an Amazon ID.

Answer: (SHOW ANSWER)

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: 91

A DevOps engineer needs to prevent any product files containing unannounced product IDs (prefixed with a specific UUID) from being stored in the production S3 bucket that users can access.

The engineer wants to use Amazon Macie to scan the production S3 bucket for unannounced product IDs. The engineer wants to use Amazon Macie to scan the production S3 bucket for unannounced product IDs. The engineer wants to use Amazon Macie to scan the production S3 bucket for unannounced product IDs.

The engineer wants to use Amazon Macie to scan the production S3 bucket for unannounced product IDs. The engineer wants to use Amazon Macie to scan the production S3 bucket for unannounced product IDs. The engineer wants to use Amazon Macie to scan the production S3 bucket for unannounced product IDs.

The engineer wants to use Amazon Macie to scan the production S3 bucket for unannounced product IDs. The engineer wants to use Amazon Macie to scan the production S3 bucket for unannounced product IDs. The engineer wants to use Amazon Macie to scan the production S3 bucket for unannounced product IDs.

A. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs.

The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs.

B. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs.

C. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs.

D. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs. The engineer should create a new S3 bucket for unannounced product IDs.

Answer: (SHOW ANSWER)

The requirement is to prevent any product files containing unannounced product IDs (prefixed with a specific UUID) from being stored in the production S3 bucket that users can access.

To achieve this, a best practice is to use a staging bucket as a control point before files go to production, combined with Amazon Macie's data classification capabilities.

- * Creating a custom data identifier in Amazon Macie allows precise detection of product IDs starting with the specific UUID, which default managed identifiers will not detect.
- * By running a Macie sensitive data discovery job on the staging bucket, you can identify files containing these sensitive product IDs.
- * Only files without findings (i.e., files that do not contain unannounced product IDs) are copied to the production bucket, ensuring no sensitive information is exposed. This approach aligns with AWS best practices for data classification and staged deployment workflows, maximizing control and reducing risk. Using Macie on the production bucket directly (options B and D) risks exposing sensitive data before detection and deletion. Option C uses managed data identifiers, which will likely not detect the custom UUID prefix pattern.

Reference from AWS Official Documentation and Study Guide:

* Amazon Macie Custom Data Identifiers: " You can create custom data identifiers in Amazon Macie to find sensitive data that is unique to your organization. " (Amazon Macie User Guide)

* Data Security Best Practices: " Use staging environments to inspect and sanitize data before moving it to production to reduce exposure risks. " (AWS Security Best Practices)

DOP-C02-KR <https://www.dumpst.com/Amazon/DOP-C02-KR-dump.html> (439 Q&As Dumps, 30%OFF Special Discount: **KrDump**)

NEW QUESTION: 92

A company is migrating its application to AWS. The application uses Amazon EC2 instances and Amazon S3 buckets. The company wants to ensure that the application can be deployed to any AWS region without manual intervention. The company is using AWS CloudFormation to create the infrastructure. The company is also using AWS CodePipeline to build and deploy the application. The company is using AWS Systems Manager Parameter Store to store configuration data. The company is using Amazon SNS to send notifications. The company is using Amazon Lambda to run serverless functions. The company is using Amazon EC2 Image Builder to create and manage AMIs. The company is using Amazon S3 to store application artifacts. The company is using Amazon S3 to store application logs. The company is using Amazon S3 to store application backups. The company is using Amazon S3 to store application data. The company is using Amazon S3 to store application metadata. The company is using Amazon S3 to store application configuration files. The company is using Amazon S3 to store application templates. The company is using Amazon S3 to store application scripts. The company is using Amazon S3 to store application images. The company is using Amazon S3 to store application videos. The company is using Amazon S3 to store application audio files. The company is using Amazon S3 to store application documents. The company is using Amazon S3 to store application spreadsheets. The company is using Amazon S3 to store application presentations. The company is using Amazon S3 to store application PDFs. The company is using Amazon S3 to store application e-books. The company is using Amazon S3 to store application music files. The company is using Amazon S3 to store application video files. The company is using Amazon S3 to store application image files. The company is using Amazon S3 to store application font files. The company is using Amazon S3 to store application CSS files. The company is using Amazon S3 to store application JavaScript files. The company is using Amazon S3 to store application HTML files. The company is using Amazon S3 to store application XML files. The company is using Amazon S3 to store application JSON files. The company is using Amazon S3 to store application YAML files. The company is using Amazon S3 to store application TOML files. The company is using Amazon S3 to store application INI files. The company is using Amazon S3 to store application properties files. The company is using Amazon S3 to store application configuration files. The company is using Amazon S3 to store application templates. The company is using Amazon S3 to store application scripts. The company is using Amazon S3 to store application images. The company is using Amazon S3 to store application videos. The company is using Amazon S3 to store application audio files. The company is using Amazon S3 to store application documents. The company is using Amazon S3 to store application spreadsheets. The company is using Amazon S3 to store application presentations. The company is using Amazon S3 to store application PDFs. The company is using Amazon S3 to store application e-books. The company is using Amazon S3 to store application music files. The company is using Amazon S3 to store application video files. The company is using Amazon S3 to store application image files. The company is using Amazon S3 to store application font files. The company is using Amazon S3 to store application CSS files. The company is using Amazon S3 to store application JavaScript files. The company is using Amazon S3 to store application HTML files. The company is using Amazon S3 to store application XML files. The company is using Amazon S3 to store application JSON files. The company is using Amazon S3 to store application YAML files. The company is using Amazon S3 to store application TOML files. The company is using Amazon S3 to store application INI files. The company is using Amazon S3 to store application properties files.

Answer: C (LEAVE A REPLY)

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/dynamic-references.html>

NEW QUESTION: 93

□ □□□□ □□ □□ Amazon EC2 □□□□□ □□□□□ □□□□□ □□□□. □□□□ □□ □□□□ □□□□□ Auto Scaling □□□□ EC2 □□□□□□ □□□□□□□□□□.

□□□□□□ Amazon S3 □□□□□ □□□□ □□□□ □□□□□ □□□□ □□□□ □□ S3 □□□□ □□□□□□. EC2 □□□□□□□□ □□ □□ □□ □□□□ □□ □□ □□ □□ □□ □□ □□□□ □□□□.

□□□□ □□ □□□□ □□□□□ □□ □□□□ □□□□□□? (2□□□ □□□□□□.)

A. S3 □□□□ □□ □□□□ □□□□ □□ IAM □□□□ □□□□□□. □□□□□ □□□□□ IAM □□□□ □□□□□□.

B. IAM □□□□□ □□□□□ □□□□□□ □□ □□□□□ □□□□□□□□□□.

C. Amazon S3□ □□ □□□□ □□□□ □□ IAM □□□□□ □□□□□□. □□ □□ □□□□ □□□□□□.

D. □□□□□ □□□□ □□□□□ □□□□□□. IAM □□□□ □□□□□ □□□□□□.

E. □□ □□□□□ □□□□□□□□□□. □ □□□ □□ □□□□ □□□□□□ □□□□ □□□□□□ □□□□□□.

Answer: (SHOW ANSWER)

To meet the requirements of deploying a workload on several hundred EC2 instances with least-privilege permissions and temporary security credentials, the company should use an IAM role and an instance profile.

An IAM role is a way to grant permissions to an entity that you trust, such as an EC2 instance. An instance profile is a container for an IAM role that you can use to pass role information to an EC2 instance when the instance starts. By using an IAM role and an instance profile, the EC2 instances can automatically receive temporary security credentials from the AWS Security Token Service (STS) and use them to access the S3 buckets. This way, the company does not need to manage or rotate any long-term credentials, such as IAM users or access keys.

To use an IAM role and an instance profile, the company should create an IAM role that has the appropriate permissions for S3 buckets. The permissions should allow the EC2 instances to read from the source S3 bucket and write to the destination S3 bucket. The company should also create a trust policy for the IAM role that specifies that EC2 is allowed to assume the role. Then, the company should add the IAM role to an instance profile. An instance profile can have only one IAM role, so the company does not need to create multiple roles or profiles for this scenario.

Next, the company should update the launch template to include the IAM instance profile. A launch template is a way to save launch parameters for EC2 instances, such as the instance type, security group, user data, and IAM instance profile. By using a launch template, the company can ensure that all EC2 instances in the Auto Scaling group have consistent configuration and permissions. The company should specify the name or ARN of the IAM instance profile in the launch template. This way, when the Auto Scaling group launches new EC2 instances based on the launch template, they will automatically receive the IAM role and its permissions through the instance profile. The other options are not correct because they do not meet the requirements or follow best practices. Creating an IAM user and generating a secret key and token is not a good option because it involves managing long-term credentials that need to be rotated regularly. Moreover, embedding credentials in user data is not secure because user data is visible to anyone who can describe the EC2 instance. Creating a trust anchor and profile is not a valid option because trust anchors are used for certificate-based authentication, not for IAM roles or instance profiles. Modifying user data to use a new secret key and token is also not a good option because it requires updating user data every time the credentials change, which is not scalable or efficient.

1: AWS Certified DevOps Engineer - Professional Certification | AWS Certification | AWS

2: DevOps Resources - Amazon Web Services (AWS)

3: Exam Readiness: AWS Certified DevOps Engineer - Professional

IAM Roles for Amazon EC2 - AWS Identity and Access Management

Working with Instance Profiles - AWS Identity and Access Management

Launching an Instance Using a Launch Template - Amazon Elastic Compute Cloud Temporary Security Credentials - AWS Identity and Access Management

NEW QUESTION: 94

DevOps uses AWS Lambda to build and deploy applications. Lambda uses Amazon Simple Queue Service (Amazon SQS) to queue messages and Amazon DynamoDB to store data. DynamoDB uses Auto Scaling to scale capacity.

DevOps uses AWS Lambda to build and deploy applications. Lambda uses Amazon Simple Queue Service (Amazon SQS) to queue messages and Amazon DynamoDB to store data. DynamoDB uses Auto Scaling to scale capacity. What is the most likely cause of the throttling issue?

- A. SQS uses the `ApproximateAgeOfOldestMessage` metric to throttle requests.
- B. SQS uses the `ApproximateAgeOfOldestMessage` metric to throttle requests. SQS uses the `NumberOfMessagesSent` metric to throttle requests.
- C. SQS uses the `NumberOfMessagesSent` metric to throttle requests. SQS uses the `NumberOfMessagesSent` metric to throttle requests.
- D. DynamoDB uses the `WriteThrottleEvents` metric to throttle requests. DynamoDB uses the `WriteThrottleEvents` metric to throttle requests.
- E. Lambda uses the `Throttles` metric to throttle requests. Lambda uses the `Throttles` metric to throttle requests.

Answer: (SHOW ANSWER)

A: If the `ApproximateAgeOfOldestMessages` indicate that orders are remaining in the SQS queue for longer than expected, the reserved concurrency limit may be set too small to keep up with the number of orders entering the queue and is being throttled. D: The DynamoDB table is using Auto Scaling. With Auto Scaling, you create a scaling policy that specifies whether you want to scale read capacity or write capacity (or both), and the minimum and maximum provisioned capacity unit settings for the table. The `ThrottledWriteRequests` metric will indicate if there is a throttling issue on the DynamoDB table, which can be resolved by increasing the maximum write capacity units for the table's Auto Scaling policy. <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

NEW QUESTION: 95

DevOps uses AWS CodeBuild to build and deploy applications. CodeBuild uses Docker to build images and Amazon S3 to store artifacts. CodeBuild uses Docker to build images and Amazon S3 to store artifacts. What is the most likely cause of the throttling issue?

- A. Docker uses Amazon Elastic Container Registry (Amazon ECR) to store images. Docker uses Amazon Elastic Container Registry (Amazon ECR) to store images.
- B. Docker uses Amazon S3 to store artifacts. Docker uses Amazon S3 to store artifacts.
- C. Docker uses Amazon Elastic Container Registry (Amazon ECR) to store images. Docker uses Amazon Elastic Container Registry (Amazon ECR) to store images.
- D. Docker uses Amazon S3 to store artifacts. Docker uses Amazon S3 to store artifacts.

Answer: A (LEAVE A REPLY)

Step 1: Storing Docker Images in Amazon ECR Docker images can be large, and storing them in a centralized, scalable location can greatly reduce build times. Amazon Elastic Container Registry (ECR) is a fully managed container registry that stores, manages, and deploys Docker container images.

Action: Store the Docker images in an ECR repository.

Why: Storing Docker images in ECR ensures that Docker images can be reused across multiple builds, improving build performance by avoiding the need to rebuild the images from scratch.

Reference: AWS documentation on Amazon ECR.

Step 2: Implementing Docker Layer Caching in CodeBuild Docker layer caching is essential for improving performance in continuous integration pipelines. CodeBuild supports local caching of Docker layers, which speeds up builds that reuse Docker images across multiple runs.

Action: Implement Docker layer caching within the CodeBuild project.

Why: This improves performance by allowing frequently used Docker layers to be cached locally, avoiding the need to pull or build the layers every time.

Reference: AWS documentation on Docker Layer Caching in CodeBuild.

This corresponds to Option A: Store the Docker images in an Amazon Elastic Container Registry (Amazon ECR) repository. Implement a local Docker layer cache for CodeBuild.

NEW QUESTION: 96

A company is using AWS Control Tower to manage multiple AWS accounts. The company wants to implement a CI/CD pipeline across all accounts. The pipeline should be able to build and deploy applications across all accounts. Which of the following is the best way to implement this?

A. Create an IAM role in the source account that has permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. The IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.

Which of the following is the best way to implement this?

- A. Create an IAM role in the source account that has permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. The IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.
- B. Create an IAM role in the source account that has permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. The IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.
- C. Create an IAM role in the source account that has permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. The IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.
- D. Create an IAM role in the source account that has permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. The IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.

Answer: A (LEAVE A REPLY)

In the source AWS account, the IAM role used by the CI/CD pipeline should have permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. The IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.

NEW QUESTION: 97

Which AWS service can be used to simulate a failure in a production environment? AWS Fault Injection Service(AWS FIS) is a service that allows you to simulate a failure in a production environment. DevOps is a set of practices that automate the deployment of an application to production. Amazon S3 is a simple object storage service. Amazon EventBridge is a serverless event-driven architecture.

A. AWS FIS is a service that allows you to simulate a failure in a production environment. AWS FIS is a service that allows you to simulate a failure in a production environment. AWS CodeBuild is a fully managed build service. CodeBuild is a fully managed build service. CodeBuild is a fully managed build service.

B. Amazon EventBridge is a serverless event-driven architecture. Amazon EventBridge is a serverless event-driven architecture. Amazon EventBridge is a serverless event-driven architecture.

C. AWS FIS is a service that allows you to simulate a failure in a production environment. AWS FIS is a service that allows you to simulate a failure in a production environment. AWS Lambda is a serverless compute service. Lambda is a serverless compute service. Lambda is a serverless compute service.

D. AWS FIS is a service that allows you to simulate a failure in a production environment. AWS FIS is a service that allows you to simulate a failure in a production environment. AWS FIS is a service that allows you to simulate a failure in a production environment.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 98

Which AWS service can be used to monitor the performance of Amazon EKS clusters? CloudWatch is a monitoring and observability service. CloudWatch is a monitoring and observability service. CloudWatch is a monitoring and observability service. DevOps is a set of practices that automate the deployment of an application to production.

Which AWS service can be used to monitor the performance of Amazon EKS clusters?

A. AWS X-Ray is a distributed tracing service. AWS X-Ray is a distributed tracing service. AWS X-Ray is a distributed tracing service. X-Ray SDK is a software development kit for X-Ray. X-Ray SDK is a software development kit for X-Ray. X-Ray SDK is a software development kit for X-Ray.

B. EKS is a managed Kubernetes service. EKS is a managed Kubernetes service. EKS is a managed Kubernetes service. CloudWatch Container Insights is a monitoring and observability service. Container Insights is a monitoring and observability service. Container Insights is a monitoring and observability service.

C. CloudWatch is a monitoring and observability service. CloudWatch is a monitoring and observability service. CloudWatch is a monitoring and observability service. SNS is a notification service. SNS is a notification service. SNS is a notification service.

D. Increasing timeouts is a way to mask the issue and does not diagnose it. Increasing timeouts is a way to mask the issue and does not diagnose it. Increasing timeouts is a way to mask the issue and does not diagnose it.

Answer: A (LEAVE A REPLY)

* AWS X-Ray provides distributed tracing, which allows visualization of latencies and errors within microservices, pinpointing bottlenecks or delays.

* Instrumenting application code with the X-Ray SDK and running the X-Ray daemon as a DaemonSet in EKS ensures tracing data is collected cluster-wide.

* Container Insights (Option B) provides resource-level metrics but not detailed request tracing.

* CloudWatch alarms and alerts (Option C) detect symptoms but don't provide root cause tracing.

* Increasing timeouts (Option D) only masks the issue and does not diagnose it.

References:

AWS X-Ray for Amazon EKS

Distributed Tracing in Microservices

NEW QUESTION: 99

Which AWS service can be used to automate the deployment of an application to production? AWS CDK(AWS Cloud Development Kit) is a framework for defining and provisioning cloud infrastructure. AWS CDK is a framework for defining and provisioning cloud infrastructure. AWS CDK is a framework for defining and provisioning cloud infrastructure. DevOps is a set of practices that automate the deployment of an application to production.

Which AWS service can be used to automate the deployment of an application to production?

B. Amazon GuardDuty is a managed service that monitors for malicious activity and potential security breaches in your AWS accounts and on-premises environments. GuardDuty continuously scans for suspicious activity and sends you notifications when it detects any potential security issues. GuardDuty is a managed service that monitors for malicious activity and potential security breaches in your AWS accounts and on-premises environments. GuardDuty continuously scans for suspicious activity and sends you notifications when it detects any potential security issues. GuardDuty is a managed service that monitors for malicious activity and potential security breaches in your AWS accounts and on-premises environments. GuardDuty continuously scans for suspicious activity and sends you notifications when it detects any potential security issues.

C. limited-ssh is a managed service that monitors for malicious activity and potential security breaches in your AWS accounts and on-premises environments. limited-ssh continuously scans for suspicious activity and sends you notifications when it detects any potential security issues. limited-ssh is a managed service that monitors for malicious activity and potential security breaches in your AWS accounts and on-premises environments. limited-ssh continuously scans for suspicious activity and sends you notifications when it detects any potential security issues.

D. Amazon Inspector is a managed service that monitors for malicious activity and potential security breaches in your AWS accounts and on-premises environments. Amazon Inspector continuously scans for suspicious activity and sends you notifications when it detects any potential security issues. Amazon Inspector is a managed service that monitors for malicious activity and potential security breaches in your AWS accounts and on-premises environments. Amazon Inspector continuously scans for suspicious activity and sends you notifications when it detects any potential security issues.

Answer: A (LEAVE A REPLY)

<https://aws.amazon.com/premiumsupport/knowledge-center/monitor-security-group-changes-ec2/>

NEW QUESTION: 102

DevOps is a culture and practice that emphasizes collaboration and communication between development and operations teams. DevOps is a culture and practice that emphasizes collaboration and communication between development and operations teams. DevOps is a culture and practice that emphasizes collaboration and communication between development and operations teams. DevOps is a culture and practice that emphasizes collaboration and communication between development and operations teams.

A. QUEUED is a CodePipeline V2 execution mode that ensures that pipeline executions run sequentially rather than replacing in-progress executions. QUEUED is a CodePipeline V2 execution mode that ensures that pipeline executions run sequentially rather than replacing in-progress executions. QUEUED is a CodePipeline V2 execution mode that ensures that pipeline executions run sequentially rather than replacing in-progress executions.

B. SUPERSEDED is a CodePipeline V2 execution mode that cancels the running execution, which violates the requirement. SUPERSEDED is a CodePipeline V2 execution mode that cancels the running execution, which violates the requirement. SUPERSEDED is a CodePipeline V2 execution mode that cancels the running execution, which violates the requirement.

C. SUPERSEDED is a CodePipeline V1 execution mode that cancels the running execution, which violates the requirement. SUPERSEDED is a CodePipeline V1 execution mode that cancels the running execution, which violates the requirement. SUPERSEDED is a CodePipeline V1 execution mode that cancels the running execution, which violates the requirement.

D. QUEUED is a CodePipeline V1 execution mode that ensures that pipeline executions run sequentially rather than replacing in-progress executions. QUEUED is a CodePipeline V1 execution mode that ensures that pipeline executions run sequentially rather than replacing in-progress executions. QUEUED is a CodePipeline V1 execution mode that ensures that pipeline executions run sequentially rather than replacing in-progress executions.

Answer: (SHOW ANSWER)

The requirements clearly indicate the need for modern CodePipeline capabilities, strict execution ordering, Git tag-based triggers, and loose coupling between pipelines. CodePipeline V2 introduces execution modes such as QUEUED and SUPERSEDED, along with native support for advanced trigger filtering when using AWS CodeConnections.

The requirement that the pipeline must wait for the previous run to finish directly maps to QUEUED mode, which ensures that pipeline executions run sequentially rather than replacing in-progress executions.

SUPERSEDED mode would cancel the running execution, which violates the requirement.

Triggering the pipeline when new Git tags are pushed is supported through CodePipeline V2 trigger filters using refs/tags/*. Branch-based triggers would not satisfy this condition.

Finally, the requirement that an existing deployment pipeline runs in response to new container images is best met using Amazon EventBridge, which natively emits events for ECR image push actions. EventBridge allows decoupled, event-driven orchestration between pipelines without tight dependencies or custom scripting. This is the AWS-recommended approach for pipeline-to-pipeline coordination.

Options C and D rely on CodePipeline V1, which lacks modern trigger filtering and execution control.

Option B incorrectly uses SUPERSEDED mode and branch-based triggers.

Therefore, Option A correctly combines CodePipeline V2, QUEUED execution mode, tag-based triggers, and EventBridge-driven pipeline chaining, meeting all requirements with best practices and minimal operational complexity.

NEW QUESTION: 103

□ □□□□ □□ □□ □□□ □□□□ □□□ □□ □□ AWS □□□□ □□□ □ □□ □□□ □□□□. □□□ AWS CodeCommit □□ □□□ □□ □□ □□□ □□□□□. DevOps □□□□□ □□□ □□ □□□□ □□□ □□□ □ □□ □□□ □□□□ □□□. □□□ □□ □□□ □□□□ □□ □□ □□□ □ □□ Git □□□ □□ □□ URL□ □□□ □ □□□□. □□ □□□□ □□□ □□ □□□ □□□□□?

- A. □□ □□□ CodeCommit □□□□□□ □□□□□. AWS CodeBuild □□□□□ □□□□ □□ □□□ CodeCommit □□□□□□ □ □□□□ CodeCommit □□□□□□□ Git □□□□□ □□□ □□□□□. CodeBuild □□□□□ □□□□□ AWS Lambda □□□ □□□□ □. □□ □□□ CodeCommit □□□□□□□□ □□ □□□□ □□□□□ Amazon EventBridge □□□ □□□□□. Lambda □□□ □□□□ □ EventBridge □□□ □□□□□.
- B. □□ □□□ Amazon S3 □□□ □□□□□. □□ □□□ CodeCommit □□□□□□ □□ Git □□ □□□ □□□□ □□□ S3 □□□ □□□□□ AWS Fargate □□□ □□□□□. Fargate □□□ □□□□□ AWS Lambda □□□ □□□□□. CodeCommit □□□□□□ □□ □□□□ □□□□□ Amazon EventBridge □□□ □□□□□. Lambda □□□ □□□□□□ EventBridge □□□ □□□□□.
- C. □□ □□□ AWS CodeArtifact □□□□□□ □□□□□. □□ □□□ □□ □□ □□□ CodeCommit □□□□□□ □□□□□ AWS CodePipeline □□□□□□ □□□□□. CodeCommit □□□□□ □□□□ □□□□□ □ □□□ CodeCommit □□□□□□ □□□ □ CodeArtifact □□□□□□ □□□□ □□□□ □□□□□□□□ □□ □□ □□□ □□□□□.
- D. □□ □□□ AWS Cloud9 □□□ CodeCommit □□□□□□ □□□□□. □□ □□□ CodeCommit □□□□□□□ AWS Cloud9 □□□ □□ □□□□□□ □□□□□. □□ □□□ CodeCommit □□□□□□□ AWS Cloud9 □□□ □□□□□.

Answer: A (LEAVE A REPLY)

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 function¹². 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 event¹.

AWS CodeCommit User Guide on resilience and disaster recovery¹.

AWS Documentation on monitoring CodeCommit events in Amazon EventBridge and Amazon CloudWatch Events².

NEW QUESTION: 104

□ □□□ □□□□□ □□□ Python □□□□ □□□□ □□ AWS CodeArtifact □□□□ □□□□□. DevOps □□□□□ AWS CodeDeploy□ □□□□ Amazon EC2 □□□□□ □□□□□□□□ □□□□ □□□. □ □□□□□□□□ CodeArtifact □□□□ □□□ Python □□□□ □□□□□. BeforeInstall □□ □□ □□□ □□□ □□□□ □□□□□. DevOps □□□□□ EC2 □□□□□ CodeArtifact □□□□ □□ □□□ □□□ □□□□ □□□. □□ □□□□ □ □□ □□□ □□□□□□□?

- A. CodeArtifact□ □□ □□□ □□ □□□ □□□□□. □ □□□□ EC2 □□□□□ □□□□□. □□□□□□□ aws codeartifact get-authorization-token CLI □□□ □□□□□.
- B. EC2 □□□□ □□□ □□ Read-FromRepository □□□ □□□□□ CodeArtifact □□□□□ □□ □□□ □□ □□□ □□□□□.
- C. EC2 □□□□□ Python □□□□□ □□□□□ □ □□□□ CodeArtifact □□□□□ ACL□ □□□□□.
- D. CodeArtifact□ □□□□ □ □□ IAM □□□ □□□ □□□□ □□□□□ □□□□□. □□□□□ □□□□□ EC2 □□□□□ □□□□ □. □□□□□□□ aws codeartifact login CLI □□□ □□□□□.

Answer: D (LEAVE A REPLY)

To allow an EC2 instance to access CodeArtifact, an IAM role attached via an instance profile must be granted permissions to access the CodeArtifact repository. The EC2 instance assumes this role.

The instance then uses the AWS CLI command `aws codeartifact login` to authenticate and configure the package manager (e.g., `pip`) to use the CodeArtifact repository. This command obtains an authorization token and sets up repository credentials securely on the instance.

Service-linked roles (Option A) are managed by AWS services, not used for instance access. CodeArtifact does not support ACLs (Option C), and resource-based policies (Option B) do not grant access to EC2 instances by principal.

This method is standard for securely managing credentials and access to CodeArtifact in automated deployments.

Reference:

AWS CodeArtifact Access Control and Authentication: " Use IAM roles attached to compute resources and the `aws codeartifact login` command to authenticate to repositories. " (AWS CodeArtifact Developer Guide)

NEW QUESTION: 105

□□□□□ □□ □□□□ □□ □□ □□□ AWS CloudFormation □□□□ □□□□ □□□□ □□□□□. □□□□ □□□ □ □□ □□□□ □□□ □ □□□□ □□□□□ □□□□□.

□□□ □□ □□□ □□□□□ □□ □□□ □□□□ □□□?

A. □□□□ CloudFormation □□□ □□□ □□□□ CloudFormation □□□ □□□ □ □□□ □□□□□. CloudFormation □□□□ □ □□ □□□□ □□□□ □□ □□□□ □□□ □□□□□.

B. □□□□ CloudFormation □□□ □□□ □□□□ CloudFormation □□□ □□□ □ □□□ □□□□□. AWS Config □□□ □□□□ □□□□ □□ □□□□ □□□ □□□ □□□□□.

C. □□□□ AWS Service Catalog □□ □□□□ CloudFormation □□□ □□□□□ □□□□□. □□ □□ □□□ □□□□□ □□□□ □. AWS Config □□□ □□□□ □□□□ □□ □□□□ □□□ □□□ □□□□□.

D. □□□□ AWS Service Catalog □□ □□□□ CloudFormation □□□ □□□□□ □□□□□. □□□ □□ □□□ □□□□□ □□□ □□. Amazon EventBridge □□□ □□□□ □□□□ □□ □□□□ □□□ □□□ □□□□□.

Answer: C (LEAVE A REPLY)

The correct answer is C. Allowing users to deploy CloudFormation stacks using AWS Service Catalog only and enforcing the use of a launch constraint is the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. AWS Service Catalog is a service that enables organizations to create and manage catalogs of IT services that are approved for use on AWS. A launch constraint is a rule that specifies the role that AWS Service Catalog assumes when launching a product. By using a launch constraint, the DevOps engineer can control the permissions that the users have when launching a product. Using AWS Config rules to detect when resources have drifted from their expected state is the best way to automate the monitoring of the resources. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config rules are custom or managed rules that AWS Config uses to evaluate whether your AWS resources comply with your desired configurations. By using AWS Config rules, the DevOps engineer can track the changes in the resources and identify any non-compliant resources.

Option A is incorrect because allowing users to deploy CloudFormation stacks using a CloudFormation service role only is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. A

CloudFormation service role is an IAM role that CloudFormation assumes to create, update, or delete the stack resources. By using a CloudFormation service role, the DevOps engineer can control the permissions that CloudFormation has when acting on the resources, but not the permissions that the users have when launching a stack. Therefore, option A does not prevent the users from launching resources that are not approved by the company. Using CloudFormation drift detection to detect when resources have drifted from their

automatically when the queue backlog grows and scales down when traffic decreases - a fully managed, cost-efficient solution (see ECS Service Auto Scaling Developer Guide).

DOP-C02-KR <https://www.dumptop.com/Amazon/DOP-C02-KR-dump.html> (439 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)

NEW QUESTION: 107

A DevOps engineer is configuring an IAM policy for an AWS Lambda function. The policy is currently configured as follows:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    }
  ]
}
```

The engineer wants to ensure that the policy only applies to EC2 instances in the NonProduction environment. Which of the following changes should be made to the policy?

- A. Add a condition to ensure that the principal making the request is an AWS Lambda function.
- B. Narrow down the resources by specifying the ARN of EC2 instances instead of allowing all resources.
- C. Add a condition to ensure that this policy only applies to EC2 instances tagged with "Environment: NonProduction".
- D. Add a condition to ensure that this policy only applies to EC2 instances tagged with "Environment: NonProduction".
- E. Add a condition to ensure that this policy only applies to EC2 instances tagged with "Environment: NonProduction".
- F. Add a condition to ensure that this policy only applies to EC2 instances tagged with "Environment: NonProduction".

Answer: A,B,D (LEAVE A REPLY)

The engineer should make the following changes to achieve a policy of least permission:

A: Add a condition to ensure that the principal making the request is an AWS Lambda function. This ensures that only Lambda functions can execute this policy.

B: Narrow down the resources by specifying the ARN of EC2 instances instead of allowing all resources. This ensures that the policy only affects EC2 instances.

D: Add a condition to ensure that this policy only applies to EC2 instances tagged with "Environment: NonProduction". This ensures that production environments are not affected by this policy.

AWS Identity and Access Management (IAM) - AWS Documentation

Certified DevOps Engineer - Professional (DOP-C02) Study Guide (page 179)

NEW QUESTION: 108

A DevOps engineer is configuring an IAM policy for an AWS CodeDeploy application. The policy is currently configured as follows:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:Describe*",
      "Resource": "*"
    }
  ]
}
```

The engineer wants to ensure that the policy only applies to EC2 instances in the NonProduction environment. Which of the following changes should be made to the policy?

A. Add a condition to ensure that the principal making the request is an AWS CodeDeploy application.

B. Narrow down the resources by specifying the ARN of EC2 instances instead of allowing all resources.

D). Configure an Amazon EventBridge rule to monitor AWS Health events. Use an Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team. This setup allows for automated monitoring and alerting of the application team in case of deployment failures or other health events⁵⁶.

AWS Elastic Beanstalk documentation on deploying Ruby applications¹.

AWS documentation on application auto-scaling⁷.

AWS documentation on automated deployment strategies with automatic rollbacks and alerts⁴⁵⁶.

NEW QUESTION: 113

You are managing a fleet of Amazon EC2 instances. You need to ensure that all instances in the fleet are up-to-date with the latest security patches. Which of the following options is the most efficient and scalable way to ensure that all instances in the fleet are up-to-date with the latest security patches?

- A. Use AWS Systems Manager Patch Manager to manage patches across the fleet. Use AWS Config to monitor compliance with patching policies.
- B. SSH into each EC2 instance and manually apply updates. This ensures that all instances are up-to-date, but it is not scalable.
- C. Use Amazon CloudWatch to monitor instance health. Use Auto Scaling to replace instances that are not healthy with new instances based on the latest AMI. This ensures that the fleet is always up-to-date.
- D. Use AWS CloudFormation to create instances with the latest AMI. Use AWS CloudTrail to monitor API activity. This ensures that all instances are created with the latest AMI.

Answer: A (LEAVE A REPLY)

Option A is the most correct because it provides both : (1) automated patching and (2) compliance monitoring/enforcement across a fleet, using AWS-native services built for exactly this purpose.

* AWS Systems Manager Patch Manager is designed to automate patching of managed instances using patch baselines , maintenance windows (or on-demand), and it produces compliance status for patching. It's the standard AWS service to apply OS/security patches at scale without SSH'ing into instances.

* AWS Config can be used to evaluate and track compliance over time against defined rules, giving centralized visibility and continuous compliance assessment. With remediation , Config can invoke Systems Manager Automation documents to correct non-compliant resources or trigger patch actions (depending on the rule/remediation design). This meets the "monitor and enforce" requirement.

Why the other options don't meet requirements as well:

* B is manual, doesn't scale well, and increases operational risk (key management, human error). It's not "automated monitoring and enforcement."

* C (replacing instances with new AMIs) can be part of an immutable infrastructure strategy, but by itself it does not provide compliance monitoring across the current fleet, and scaling policies based on CloudWatch metrics are unrelated to patch compliance. Also, patch cadence would depend on AMI pipelines and instance rotation rather than direct compliance enforcement.

* D is operationally heavy and mismatched: CloudTrail records API activity; it does not natively provide "patch compliance" status for instance OS packages. Recreating instances via CloudFormation for every patch is not an efficient or standard enforcement mechanism for patch compliance.

NEW QUESTION: 114

You are managing a fleet of Amazon EC2 instances across multiple AWS Organizations. Which of the following options is the most efficient and scalable way to ensure that all instances in the fleet are up-to-date with the latest security patches?

* Amazon Linux AMI is pre-installed with Chef.

* Amazon Linux AMI is pre-installed with Chef.

* Amazon Linux AMI is pre-installed with Chef.

Which of the following is the correct answer?

A. Amazon EC2 Image Builder is a service that automates the creation, management, and deployment of customized, secure, and up-to-date server images that are pre-installed with software and configuration settings tailored to meet specific IT standards. EC2 Image Builder simplifies the creation and maintenance of golden images, and makes it easy to generate images for multiple platforms, such as Amazon EC2 and on-premises. EC2 Image Builder also integrates with AWS Resource Access Manager, which allows you to share your images across accounts within your organization or with external AWS accounts. This solution meets the requirements of automating the tasks of updating the Linux AMIs, installing the Chef agent, and providing the images to the department's accounts with the least management overhead. References:

B. Amazon EC2 Image Builder is a service that automates the creation, management, and deployment of customized, secure, and up-to-date server images that are pre-installed with software and configuration settings tailored to meet specific IT standards. EC2 Image Builder simplifies the creation and maintenance of golden images, and makes it easy to generate images for multiple platforms, such as Amazon EC2 and on-premises. EC2 Image Builder also integrates with AWS Resource Access Manager, which allows you to share your images across accounts within your organization or with external AWS accounts. This solution meets the requirements of automating the tasks of updating the Linux AMIs, installing the Chef agent, and providing the images to the department's accounts with the least management overhead. References:

C. AWS Systems Manager Automation Runbook is a service that automates the creation, management, and deployment of customized, secure, and up-to-date server images that are pre-installed with software and configuration settings tailored to meet specific IT standards. EC2 Image Builder simplifies the creation and maintenance of golden images, and makes it easy to generate images for multiple platforms, such as Amazon EC2 and on-premises. EC2 Image Builder also integrates with AWS Resource Access Manager, which allows you to share your images across accounts within your organization or with external AWS accounts. This solution meets the requirements of automating the tasks of updating the Linux AMIs, installing the Chef agent, and providing the images to the department's accounts with the least management overhead. References:

D. Amazon EC2 Image Builder is a service that automates the creation, management, and deployment of customized, secure, and up-to-date server images that are pre-installed with software and configuration settings tailored to meet specific IT standards. EC2 Image Builder simplifies the creation and maintenance of golden images, and makes it easy to generate images for multiple platforms, such as Amazon EC2 and on-premises. EC2 Image Builder also integrates with AWS Resource Access Manager, which allows you to share your images across accounts within your organization or with external AWS accounts. This solution meets the requirements of automating the tasks of updating the Linux AMIs, installing the Chef agent, and providing the images to the department's accounts with the least management overhead. References:

Answer: B (LEAVE A REPLY)

Amazon EC2 Image Builder is a service that automates the creation, management, and deployment of customized, secure, and up-to-date server images that are pre-installed with software and configuration settings tailored to meet specific IT standards. EC2 Image Builder simplifies the creation and maintenance of golden images, and makes it easy to generate images for multiple platforms, such as Amazon EC2 and on-premises. EC2 Image Builder also integrates with AWS Resource Access Manager, which allows you to share your images across accounts within your organization or with external AWS accounts. This solution meets the requirements of automating the tasks of updating the Linux AMIs, installing the Chef agent, and providing the images to the department's accounts with the least management overhead. References:

* Amazon EC2 Image Builder

* Sharing EC2 Image Builder images

NEW QUESTION: 115

DevOps is a culture and practice that emphasizes collaboration and communication between development and operations teams. Which of the following is the correct answer?

Amazon CloudWatch Logs is a service that provides a central location for monitoring and logging your AWS resources. Amazon CloudWatch Logs allows you to monitor and log your AWS resources in real-time. Amazon CloudWatch Logs also provides a central location for monitoring and logging your AWS resources. Amazon CloudWatch Logs also provides a central location for monitoring and logging your AWS resources.

DevOps is a culture and practice that emphasizes collaboration and communication between development and operations teams. Which of the following is the correct answer?

A. Amazon CloudWatch Synthetics is a service that provides a central location for monitoring and logging your AWS resources. Amazon CloudWatch Synthetics allows you to monitor and log your AWS resources in real-time. Amazon CloudWatch Synthetics also provides a central location for monitoring and logging your AWS resources. Amazon CloudWatch Synthetics also provides a central location for monitoring and logging your AWS resources.

B. Amazon CloudWatch Synthetics is a service that provides a central location for monitoring and logging your AWS resources. Amazon CloudWatch Synthetics allows you to monitor and log your AWS resources in real-time. Amazon CloudWatch Synthetics also provides a central location for monitoring and logging your AWS resources. Amazon CloudWatch Synthetics also provides a central location for monitoring and logging your AWS resources.

C. Amazon GuardDuty is a service that provides a central location for monitoring and logging your AWS resources. Amazon GuardDuty allows you to monitor and log your AWS resources in real-time. Amazon GuardDuty also provides a central location for monitoring and logging your AWS resources. Amazon GuardDuty also provides a central location for monitoring and logging your AWS resources.

D. AWS Firewall Manager sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

Answer: (SHOW ANSWER)

"The firewall appliance sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO"

NEW QUESTION: 116

AWS Organizations sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

AWS CloudFormation sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

AWS CodeConnections sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

AWS CodePipeline sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO. (300 points)

A. AWS Config sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

B. AWS Config sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

C. CodeConnections sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

D. CodeConnections sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

E. IAM sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

F. IAM sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO.

Answer: A,C,E (LEAVE A REPLY)

* Step A: Using a tool like CloudFormation resource import or IaC generator to scan and create a template from existing resources is efficient to bring current infrastructure under management.

* Step C: Using CodeConnections (AWS 's solution to connect Git repositories) with AWS CodePipeline ensures any changes to CloudFormation templates in the Git repo automatically deploy infrastructure changes, enforcing infrastructure as code workflows.

* Step E: Creating an IAM role with CloudFormation as the principal ensures CloudFormation has permissions to manage resources. Using an SCP to deny all actions except by this role enforces strict control, preventing manual changes outside the pipeline. Option B uses AWS Config which is more for compliance and auditing, not direct resource import. Option D is invalid because CloudFormation does not natively sync with Git; CodePipeline does. Option F is less secure than denying all except the IAM role.

Reference:

AWS CloudFormation Resource Import: " Import existing resources into CloudFormation stacks for management. " (CloudFormation Resource Import) AWS CodePipeline and CodeConnections Integration: " Use CodeConnections to connect Git providers with AWS CodePipeline for continuous deployment. " (AWS CodePipeline Git Integration) AWS Organizations SCP and IAM Role Best Practices: " Use SCPs to restrict actions and IAM roles with limited principals to enforce secure management. " (AWS Organizations Best Practices)

NEW QUESTION: 117

Which of the following AWS DevOps tools can be used to monitor and alert on AWS Config rule violations? (Select TWO.)

Options:

A. Amazon S3, AWS Config, Amazon Athena, Amazon GuardDuty, Amazon CloudWatch

B. Amazon CloudWatch, Amazon GuardDuty, Amazon CloudTrail

C. Amazon SNS, Amazon EventBridge, AWS Management Console, Amazon CloudWatch, Amazon EventBridge

D. Amazon SNS, Amazon EventBridge, AWS Lambda, AWS CloudTrail, AWS API Gateway, Amazon EventBridge

Answer: D (LEAVE A REPLY)

Create an Amazon EventBridge Rule Using an AWS CloudTrail Event Pattern:

AWS CloudTrail logs API calls made in your account, including actions performed by roles.

Create an EventBridge rule that matches CloudTrail events where the AssumeRole API call is made to assume the administrator role.

Invoke an AWS Lambda Function:

Configure the EventBridge rule to trigger a Lambda function whenever the rule's conditions are met.

The Lambda function will handle the logic to send a notification.

Publish a Message to an Amazon SNS Topic:

The Lambda function will publish a message to an SNS topic to notify the security team.

Subscribe the security team's email address to this SNS topic to receive real-time notifications.

Example EventBridge rule pattern:

```
{
  "source": ["aws.cloudtrail"],
  "detail-type": ["AWS API Call via CloudTrail"],
  "detail": {
    "eventSource": ["sts.amazonaws.com"],
    "eventName": ["AssumeRole"],
    "requestParameters": {
      "roleArn": ["arn:aws:iam::<account-id>:role/AdministratorRole"]
    }
  }
}
```

Example Lambda function (Node.js) to publish to SNS:

```
const AWS = require('aws-sdk');
const sns = new AWS.SNS();
exports.handler = async (event) => {
  const params = {
    Message: `Administrator role assumed: ${JSON.stringify(event.detail)}`, TopicArn: 'arn:aws:sns:<region>:<account-id>:<sns-topic>'
  };
  await sns.publish(params).promise();
};
```

};

References:

Creating EventBridge Rules

Using AWS Lambda with Amazon SNS

NEW QUESTION: 118

A company is migrating its on-premises data center to AWS. The company has 400 AWS accounts and is using a multi-account strategy. The company is using a multi-account strategy to manage its AWS accounts. The company is using a multi-account strategy to manage its AWS accounts. The company is using a multi-account strategy to manage its AWS accounts.

The company is using Linux-based servers and is using NetApp ONTAP storage. The company is using Linux-based servers and is using NetApp ONTAP storage. The company is using Linux-based servers and is using NetApp ONTAP storage. The company is using Linux-based servers and is using NetApp ONTAP storage.

DevOps teams are using AWS CloudFormation to manage their infrastructure. DevOps teams are using AWS CloudFormation to manage their infrastructure. DevOps teams are using AWS CloudFormation to manage their infrastructure. DevOps teams are using AWS CloudFormation to manage their infrastructure.

Which of the following is the best solution to replicate data from the on-premises storage to AWS?

- A. Use Amazon S3 as the target storage. Use S3 as the target storage. Use S3 as the target storage. Use S3 as the target storage.
- B. Use Amazon FSx for Windows as the target storage. Use Amazon FSx for Windows as the target storage. Use Amazon FSx for Windows as the target storage. Use Amazon FSx for Windows as the target storage.
- C. Use Amazon FSx for NetApp ONTAP as the target storage. Use Amazon FSx for NetApp ONTAP as the target storage. Use Amazon FSx for NetApp ONTAP as the target storage. Use Amazon FSx for NetApp ONTAP as the target storage.
- D. Use Amazon Elastic File System (Amazon EFS) as the target storage. Use Amazon Elastic File System (Amazon EFS) as the target storage. Use Amazon Elastic File System (Amazon EFS) as the target storage. Use Amazon Elastic File System (Amazon EFS) as the target storage.

Answer: C (LEAVE A REPLY)

Amazon FSx for NetApp ONTAP provides NetApp ONTAP features in AWS, including SnapMirror replication and storage efficiencies like deduplication and compression. Create FSx for ONTAP in each Region and use SnapMirror from on-prem ONTAP to each Region for efficient, incremental replication.

Regions can serve data independently of on-prem availability once replicated.

NEW QUESTION: 119

A company is using Amazon EKS to run its Kubernetes clusters. The company is using Amazon EKS to run its Kubernetes clusters. The company is using Amazon EKS to run its Kubernetes clusters. The company is using Amazon EKS to run its Kubernetes clusters.

Which of the following is the best solution to monitor the health of the EKS clusters?

- A. Amazon CloudWatch Logs to monitor EKS logs. Amazon CloudWatch Logs to monitor EKS logs. Amazon CloudWatch Logs to monitor EKS logs. Amazon CloudWatch Logs to monitor EKS logs.
- B. Amazon CloudWatch Logs to monitor EKS logs. Amazon CloudWatch Logs to monitor EKS logs. Amazon CloudWatch Logs to monitor EKS logs. Amazon CloudWatch Logs to monitor EKS logs.
- C. EKS logs to monitor Amazon S3 logs. EKS logs to monitor Amazon S3 logs. EKS logs to monitor Amazon S3 logs. EKS logs to monitor Amazon S3 logs.
- D. EKS logs to monitor Amazon S3 logs. EKS logs to monitor Amazon S3 logs. EKS logs to monitor Amazon S3 logs. EKS logs to monitor Amazon S3 logs.

Answer: A (LEAVE A REPLY)

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html#LambdaFunctionExample>

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html>

NEW QUESTION: 120

DevOps wants to use AWS CloudFormation to create Amazon ECS instances on Amazon EC2 Auto Scaling. The CloudFormation template should create an ECS cluster and EC2 instances that register with the ECS cluster.

Which of the following is the correct way to create the ECS cluster and EC2 instances?

A. AWS: ECS: Cluster resource to create the ECS cluster and AWS: ECS: InstanceProfile resource to create the EC2 instances.

B. UserData property of the AWS: AutoScaling: LaunchConfiguration resource to create the EC2 instances and ECS: Cluster resource to create the ECS cluster.

C. AWS:EC2: Instance resource to create the EC2 instances and ECS: Cluster resource to create the ECS cluster.

D. AWS: CloudFormation: CustomResource resource to create the ECS cluster and AWS Lambda function to create the EC2 instances.

Answer: B (LEAVE A REPLY)

The UserData property of the AWS: AutoScaling: LaunchConfiguration resource can be used to specify a script that runs when the EC2 instances are launched. This script can include the ECS cluster name as an environment variable for the ECS agent running on the EC2 instances. This way, the EC2 instances will register with the correct ECS cluster. Option A is incorrect because the AWS: ECS: Cluster resource does not have a property to reference the EC2 instances. Option C is incorrect because the EC2 instances are launched by the Auto Scaling group, not by the AWS: EC2: Instance resource. Option D is incorrect because using a custom resource and a Lambda function is unnecessary and overly complex for this scenario. References: AWS::AutoScaling::LaunchConfiguration, Amazon ECS Container Agent Configuration

NEW QUESTION: 121

A company wants to migrate its Microsoft SQL Server database to Amazon RDS. The database is currently on-premises and the company requires RPO < 1 min and RTO < 10 min. Which of the following is the best migration strategy?

Which of the following is the best migration strategy?

A. RDS Multi-AZ DB instance with Read Replica in another Region. Route 53 failover redirects application traffic automatically.

B. RDS Multi-AZ DB instance with Read Replica in another Region. Route 53 failover redirects application traffic automatically.

C. RDS Multi-AZ DB instance with Read Replica in another Region. Route 53 failover redirects application traffic automatically.

D. RDS Multi-AZ DB instance with Read Replica in another Region. Route 53 failover redirects application traffic automatically.

Answer: (SHOW ANSWER)

RDS Multi-AZ cluster deployments with cross-Region read replicas support near real-time replication (<1 min RPO) and fast promotion (<10 min RTO). Route 53 CNAME failover redirects application traffic automatically. This is the AWS-recommended DR pattern.

DOP-C02-KR is a question ID. The correct answer is **DOP-C02-KR**. The question is about RDS Multi-AZ cluster deployments with cross-Region read replicas. The correct answer is D.

Special Discount: **KrDump**)

NEW QUESTION: 122

0 0000 0000 000 00 00(DR) 000 000 AWS 000 0000 000 0000000. 00 0000 0000000 Amazon EC2 0000000 0000 AWS CloudFormation 0 00 000000. 00000000 Amazon FSx for NetApp ONTAP 000 0000000 000000 000000. EC2 0000000 0000000 00000 00000 00000. DevOps 000000 0000 AMI 0000 DR 0000 000000. 00 DevOps 000000 0000 000000 000000 CloudFormation 0000 0000000000. DR 00 0 000000 1000 RPO 0000 0000. 0000 00 0000 00000 00000 0000000?

- A. 0 00 0000 Amazon S3 0000 000000. S3 0000 00 S3 00 00 00(CRR)0 000000. FSx for ONTAP 0000 0 00 00 00000 0000 S3 00000 00000 0000 AWS Lambda 0000 0000000.
- B. DR 0000 FSx for ONTAP 0000000 0000000. 00 00 NetApp SnapMirror 0 00000 0000 0000 DR 00000 000000 50 0000 0000 0000000.
- C. EC2 0000000 0000 000000 0000 0000 000000 000000 AWS Lambda 0000 0000000. 000000 DR 0000 00000 0 000000 0000000 Lambda 0000 0000000. 100000 Lambda 0000 000000 Amazon EventBridge 00 0000 0000000.
- D. AWS Backup 0 000000 00 0000 100 0000 0000 00 00 0000 0000000. 00 000000 DR 0000 0000000. 0000 0 0000 EC2 0000000 00 0000 0000000.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 123

0 0000 RDS DB 0000000 00 Amazon RDS 000000 00 0000 000000000. DevOps 00 Amazon CloudWatch 00000000 00 00 000000 00000000 0000. 00 000000 0 00 0000 0000000?

- A. RDS 0000000 000000 RDS 000000 00 00 00 000000 000000 Amazon EventBridge 0000 000000. CloudWatch 0000 00 000000 000000 AWS Lambda 0000 000000. Lambda 0000 0000000 EventBridge 0000 0000000. CloudWatch 00000000 000000 0000 00 000000 0000000000.
- B. 00 000000 0000 AWS CloudTrail 0 000000 000000 000000. 00 000000 Amazon CloudWatch Logs 0 000000 000000 00000000. RDS 000000 00 00 000000 00000000 CloudWatch Logs 00 0000 0000 000000. CloudWatch 0000000 00 00 0000 0000 0000000000.
- C. RDS 000000 00 00 0000(RDS 0000)0 000000 Amazon EventBridge 0000 000000. CloudWatch 0000 000000. CloudWatch 0000 0000 00000000 EventBridge 0000 00000000. CloudWatch 00000000 000000 00 0000 0000000000.
- D. 0000 000000 0000 AWS CloudTrail 0 000000 000000 000000. 000000 000000 0000 000000 Amazon CloudWatch Logs 0 000000. RDS 000000 00 00 000000 00000000 CloudWatch Logs 00 0000 0000 000000. CloudWatch 0000000 000000 0000 0000 0000000000.

Answer: A ([LEAVE A REPLY](#))

Step 1: Reacting to RDS Storage Autoscaling Events Using Amazon EventBridge Amazon RDS emits events when storage autoscaling occurs. To visualize these events in a CloudWatch dashboard, you can create an EventBridge rule that listens for these specific autoscaling events.

Action: Create an EventBridge rule that reacts to RDS storage autoscaling events from the RDS event stream.

Why: EventBridge allows you to listen to RDS events and route them to specific AWS services for processing.

Step 2: Creating a Custom CloudWatch Metric via Lambda Once the EventBridge rule detects a storage autoscaling event, you can use a Lambda function to publish a custom metric to CloudWatch. This metric can then be visualized in a CloudWatch dashboard.

Action: Use a Lambda function to publish custom metrics to CloudWatch based on the RDS storage autoscaling events.

Why: Custom metrics allow you to track specific events like autoscaling and visualize them easily on a CloudWatch dashboard.

Reference: AWS documentation on Publishing Custom Metrics to CloudWatch.

This corresponds to Option A: Create an Amazon EventBridge rule that reacts to RDS storage autoscaling events from RDS events.

Create an AWS Lambda function that publishes a CloudWatch custom metric.

Configure the EventBridge rule to invoke the Lambda function. Visualize the custom metric by using the CloudWatch dashboard.

NEW QUESTION: 124

- Amazon Redshift supports audit logging which can be configured to log user activity, including queries and user changes. The logs can be delivered to an S3 bucket (Option B), which is the standard location for Redshift audit logs.
- While Redshift can send some logs to CloudWatch, the native audit logging is typically routed to S3. Creating a CloudWatch dashboard with a log widget (Option D) allows visualization of user activities from logs ingested into CloudWatch Logs or through custom processing. Option A is invalid because CloudTrail does not capture Redshift query logs directly. Option C incorrectly assumes Redshift audit logs can be delivered directly to CloudWatch. Option E adds complexity by requiring Lambda and Athena unnecessarily when direct visualization in CloudWatch is simpler.
- References:
- Amazon Redshift Audit Logging
 - Visualizing Logs with Amazon CloudWatch Dashboards
- NEW QUESTION: 125**
- Amazon EC2 instances can be configured to send logs to CloudWatch. The logs can be delivered to an S3 bucket (Option B), which is the standard location for EC2 audit logs.
- While EC2 can send some logs to CloudWatch, the native audit logging is typically routed to S3. Creating a CloudWatch dashboard with a log widget (Option D) allows visualization of user activities from logs ingested into CloudWatch Logs or through custom processing. Option A is invalid because CloudTrail does not capture EC2 query logs directly. Option C incorrectly assumes EC2 audit logs can be delivered directly to CloudWatch. Option E adds complexity by requiring Lambda and Athena unnecessarily when direct visualization in CloudWatch is simpler.
- References:
- Amazon EC2 Audit Logging
 - Visualizing Logs with Amazon CloudWatch Dashboards
- NEW QUESTION: 126**
- Amazon EC2 instances can be configured to send logs to CloudWatch. The logs can be delivered to an S3 bucket (Option B), which is the standard location for EC2 audit logs.
- While EC2 can send some logs to CloudWatch, the native audit logging is typically routed to S3. Creating a CloudWatch dashboard with a log widget (Option D) allows visualization of user activities from logs ingested into CloudWatch Logs or through custom processing. Option A is invalid because CloudTrail does not capture EC2 query logs directly. Option C incorrectly assumes EC2 audit logs can be delivered directly to CloudWatch. Option E adds complexity by requiring Lambda and Athena unnecessarily when direct visualization in CloudWatch is simpler.
- References:
- Amazon EC2 Audit Logging
 - Visualizing Logs with Amazon CloudWatch Dashboards

Answer: (SHOW ANSWER)

NEW QUESTION: 125

C. Lambda EventBridge SNS
 Lambda EventBridge SNS

D. Lambda EventBridge SES
 Lambda EventBridge SES

Answer: (SHOW ANSWER)

* Systems Manager State Manager with Automation documents allows running scripts sequentially and reliably with built-in retry and status tracking.

* Using EventBridge with SNS for notifications leverages managed services with minimal custom development.

* Using SES (Options B, D) requires more setup and custom logic for email formatting and sending.

* Lambda (Options C, D) can run scripts but might have limitations on execution time and complexity compared to Systems Manager Automation.

References:

AWS Systems Manager Automation

Monitoring Automation Execution and Notifications

NEW QUESTION: 126

CloudWatch Logs 30, 90, 180 days? 30, 90, 180 days?

- A. Kinesis Data Streams S3
- B. Kinesis Data Firehose S3 One Zone- IA Glacier Flexible Retrieval
- C. Kinesis Data Streams S3 Standard-IA Glacier Instant Retrieval
- D. Kinesis Data Firehose S3 Standard-IA(90) Glacier Deep Archive(180)

Answer: D (LEAVE A REPLY)

Exporting logs using CloudWatch subscription filters + Firehose # S3 enables long-term archival. S3 lifecycle transitions optimize cost with Standard-IA (low-latency) and Deep Archive (long-term retention). This matches AWS best practice for cost-effective log retention.

NEW QUESTION: 127

AWS Organizations AWS Lambda VPC Lambda AWS Amazon Elastic File System(Amazon EFS)

Lambda Amazon EFS DevOps Lambda EPS EFS

DevOps (3)

- A. EFS B A EFS
- B. Amazon EFS SCP

C. B EFS AWS DMS(AWS Database Migration Service) A B

D. VPC EFS Lambda

E. VPC A B

F. A IAM B Lambda

Answer: A,E,F (LEAVE A REPLY)

A Lambda function in one account can mount a file system in a different account. For this scenario, you configure VPC peering between the function VPC and the file system VPC. https://docs.aws.amazon.com

/lambda/latest/dg/services-efs.html

https://aws.amazon.com/ru/blogs/storage/mount-amazon-efs-file-systems-cross-account-from-amazon-eks/

1. Need to update the file system policy on EFS to allow mounting the file system into Account B.

File System Policy

\$ cat file-system-policy.json

```
{
"Statement": [
{
"Effect": "Allow",
"Action": [
"elasticfilesystem:ClientMount",
"elasticfilesystem:ClientWrite"
],
"Principal": {
"AWS": "arn:aws:iam::<aws-account-id-A>:root" # Replace with AWS account ID of EKS cluster
}
}
]
}
```

2. Need VPC peering between Account A and Account B as the pre-requisite

3. Need to assume cross-account IAM role to describe the mounts so that a specific mount can be chosen.

NEW QUESTION: 128

DevOps AWS CloudFormation ALB(Application Load Balancer) Amazon EC2 DevOps IPv6

DevOps IPv6 CloudFormation

A. VPC IPv6 CIDR EC2 IPv6

B. EC2 IPv6 IP EC2 ALB 443

C. ALB NLB(Network Load Balancer) IPv6 CIDR VPC NLB IPv6 IP

D. IPv6 CIDR block VPC ALB subnets. Port 443 ALB listener IP address type EC2 instances. ALB targets.

Answer: D (LEAVE A REPLY)

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: 129

Amazon EC2 instances 24-hour Amazon CloudWatch Logs Auto Scaling group. What is the best way to capture logs from EC2 instances?

A. AWS Step Functions CloudWatch Logs EC2 instances AWS Lambda Amazon EventBridge

B. Amazon CloudWatch Amazon Simple Notification Service(Amazon SNS) EC2 instances

C. Amazon CloudWatch Amazon Simple Queue Service(Amazon SQS) Amazon EventBridge

D. AWS Lambda CloudWatch Logs EC2 instances Lambda Amazon EventBridge

Answer: (SHOW ANSWER)

"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>

[amazon.com/AmazonCloudWatch/latest/logs/Subscriptions.html](https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/Subscriptions.html)

NEW QUESTION: 130

AWS Lambda Amazon S3 Amazon DynamoDB. How can you store Lambda function code in S3 and have it automatically update Lambda functions when code is updated in S3?

A. Lambda S3 Lambda OnFailure Lambda S3 Lambda

B. Amazon S3 S3 Lambda Lambda Lambda Lambda

C. Amazon Simple Queue Service(Amazon SQS) Lambda OnFailure Lambda

D. Lambda /tmp S3 Lambda Lambda Lambda Lambda

Answer: B (LEAVE A REPLY)

Option A is incorrect because increasing the memory of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Increasing the memory of the Lambda function might improve its performance or reduce its execution time, but it does not affect its invocation. Moreover, increasing the memory of the Lambda function might incur higher costs, as Lambda charges based on the amount of memory allocated to the function.

Option B is correct because creating a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket is a necessary step to configure an S3 event source. A resource policy is a JSON document that defines who can access a Lambda resource and under what conditions. By granting Amazon S3 permission to invoke the Lambda function, the company ensures that the Lambda function runs when a new object is created or an existing object is modified in the S3 bucket¹.

Option C is incorrect because configuring an Amazon Simple Queue Service (Amazon SQS) queue as an On-Failure destination for the Lambda function does not help with triggering the Lambda function. An On-Failure destination is a feature that allows Lambda to send events to another service, such as SQS or Amazon Simple Notification Service (Amazon SNS), when a function invocation fails. However, this feature only applies to asynchronous invocations, and S3 event sources use synchronous invocations. Therefore, configuring an SQS queue as an On-Failure destination would have no effect on the problem.

Option D is incorrect because provisioning space in the /tmp folder of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source.

Provisioning space in the /tmp folder of the Lambda function might help with processing large files from the S3 bucket, as it provides temporary storage for up to 512 MB of data. However, it does not affect the invocation of the Lambda function.

References:

Using AWS Lambda with Amazon S3

Lambda resource access permissions

AWS Lambda destinations

[AWS Lambda file system]

NEW QUESTION: 131

A company is using GitHub Actions to build and deploy an application. The application is built using Docker and is deployed to Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. The application is currently deployed to Amazon EC2. The company wants to migrate the application to Amazon ECS. Which of the following is the most appropriate action to take?

What is the most appropriate action to take?

A. Amazon ECS on AWS CodePipeline on Amazon EC2 on Lambda.

B. AWS CodeDeploy on AWS CodePipeline on Amazon EC2 on Lambda.

C. AWS Elastic Beanstalk on AWS CodePipeline on Amazon EC2 on Lambda.

D. GitHub Actions on AWS CodeDeploy on Amazon EC2 on Lambda.

Answer: (SHOW ANSWER)

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

NEW QUESTION: 132

A company is using Amazon ECR to store Docker images. The company wants to migrate the application to Amazon ECS. Which of the following is the most appropriate action to take?

What is the most appropriate action to take?

A. S3 on Amazon ECS (on Amazon EC2).

B. Amazon ECS on Amazon ECR on Amazon EC2 on Lambda.

- C. Lambda
- D.

Answer: B (LEAVE A REPLY)

ECR Lifecycle Policies automatically manage image retention based on tag status, age, or count. They execute natively within the ECR service, requiring no external management or scripts - the least overhead and AWS-recommended cleanup method.

NEW QUESTION: 133

- A. S3 WebsiteConfiguration
- B. S3 RequestType Delete
- C. RemoveOnDeletion DeletionPolicy
- D. S3 CloudFormation S3 Empty DeletionPolicy

Answer: (SHOW ANSWER)

Step 1: Understanding the Deletion Failure
The most likely reason why the CloudFormation stack failed to delete is that the S3 bucket was not empty. AWS CloudFormation cannot delete an S3 bucket that contains objects, so if the website files are still in the bucket, the deletion will fail.

Issue: The S3 bucket is not empty during deletion, preventing the stack from being deleted.

Step 2: Modifying the Custom Resource to Handle Deletion
To mitigate this issue, you can modify the Lambda function associated with the custom resource to automatically empty the S3 bucket when the stack is being deleted. By adding logic to handle the RequestType: Delete event, the function can recursively delete all objects in the bucket before allowing the stack to be deleted.

Action: Modify the Lambda function to recursively delete the objects in the S3 bucket when RequestType is set to Delete.

Why: This ensures that the S3 bucket is empty before CloudFormation tries to delete it, preventing the stack deletion failure.

Reference: AWS documentation on CloudFormation custom resources.

This corresponds to Option B: Deletion has failed because the S3 bucket is not empty. Modify the custom resource's AWS Lambda function code to recursively empty the bucket when RequestType is Delete.

NEW QUESTION: 134

- DevOps Amazon Elastic Kubernetes Service(Amazon EKS)
- DevOps CPU Pod
- Kubernetes Metrics Server

- A.** EKS automatically scales the number of nodes based on node CPU utilization. This is done by the Cluster Autoscaler. CPU utilization is the primary metric used for scaling.
- B.** Kubernetes Horizontal Pod Autoscaler (HPA) scales the number of pods based on CPU utilization. Kubernetes Vertical Pod Autoscaler (VPA) scales the number of pods based on CPU utilization. HPA and VPA are used to scale pods.
- C.** AWS Systems Manager AWS-UpdateEKSManagedNodeGroup Automation updates the NodeGroupDesiredSize, NodeGroupMaxSize, and NodeGroupMinSize. This is done by the Cluster Autoscaler.
- D.** Kubernetes Horizontal Pod Autoscaler (HPA) scales the number of pods based on CPU utilization. Kubernetes Vertical Pod Autoscaler (VPA) scales the number of pods based on CPU utilization. HPA and VPA are used to scale pods.

Answer: D (LEAVE A REPLY)

To scale microservice Pods based on CPU utilization, the Kubernetes Horizontal Pod Autoscaler (HPA) uses the Kubernetes Metrics Server to monitor resource usage and automatically adjusts the number of Pods.

However, scaling Pods may require additional nodes if the current node capacity is insufficient.

- * The Cluster Autoscaler works with EKS managed node groups to add or remove worker nodes based on pending Pod requirements and resource usage.
- * By deploying both HPA and Cluster Autoscaler, the system can automatically scale Pods and add nodes as necessary, ensuring efficient resource utilization and availability.
- * Configuring the Cluster Autoscaler with auto-discovery allows it to manage node groups without manual intervention, reducing operational effort.
- * Option A only scales nodes based on node CPU utilization, not Pods.
- * Option B uses VPA recommender mode, which only suggests resource changes and does not scale automatically.
- * Option C involves manual updates and is not automated scaling. Therefore, option D provides the most operationally efficient, fully automated scaling solution.

Reference from AWS Official Documentation:

- * Kubernetes Horizontal Pod Autoscaler: " HPA automatically scales the number of Pods based on observed CPU utilization or other metrics. " (Kubernetes HPA)
- * Cluster Autoscaler on Amazon EKS: " The Cluster Autoscaler automatically adjusts the size of the Kubernetes cluster when there are Pods that fail to run due to insufficient resources or when nodes in the cluster are underutilized. " (AWS EKS Cluster Autoscaler)

NEW QUESTION: 135

- A.** AWS Organizations automatically scales the number of nodes based on node CPU utilization. This is done by the Cluster Autoscaler. CPU utilization is the primary metric used for scaling.
- B.** AWS Config automatically scales the number of nodes based on node CPU utilization. This is done by the Cluster Autoscaler. CPU utilization is the primary metric used for scaling.

B. CloudFormation `AWSSSMAssociation` `AWS-JoinDirectoryServiceDomain` Runbook `EC2` `AmazonSSMManagedInstanceCore` `AmazonSSMDirectoryServiceAccess` AWS IAM `AmazonSSMManagedInstanceCore` `AmazonSSMDirectoryServiceAccess` AWS IAM.

C. `AWS Managed Microsoft AD` `AWS Secrets Manager` `CloudFormation` `AWSSSMAssociation` `AWS-CreateManagedWindowsInstanceWithApproval` Runbook `EC2 Auto Scaling` `Secrets Manager` `ARN` `EC2` `IAM` `AmazonSSMDirectoryServiceAccess` `SecretsManagerReadWrite` AWS IAM.

D. `AWS Managed Microsoft AD` `AWS Secrets Manager` `CloudFormation` `EC2` `Secrets Manager` `AWS Managed Microsoft AD` `EC2` `IAM` `AmazonSSMManagedInstanceCore` `SecretsManagerReadWrite` AWS IAM.

Answer: B (LEAVE A REPLY)

To meet the requirements, the DevOps engineer needs to create a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory with the most operational efficiency. The DevOps engineer can use AWS Systems Manager Automation to automate the domain join process using an existing runbook called `AWS-JoinDirectoryServiceDomain`. This runbook can join Windows instances to an AWS Managed Microsoft AD or Simple AD directory by using PowerShell commands. The DevOps engineer can create an `AWSSSMAssociation` resource in the CloudFormation template to associate the runbook with the EC2 instances that have specific tags. The tags can be defined in the launch template and propagated on launch to the EC2 instances. The DevOps engineer can also define the required parameters for the runbook, such as the directory ID, directory name, and organizational unit. The DevOps engineer can attach the `AmazonSSMManagedInstanceCore` and `AmazonSSMDirectoryServiceAccess` AWS managed policies to the IAM role that the EC2 instances use. These policies grant the necessary permissions for Systems Manager and Directory Service operations.

NEW QUESTION: 138

The DevOps engineer needs to create a solution that uses Amazon Elastic Container Service (Amazon ECS) to run a container in the `us-west-2` region. The solution must use Amazon CloudWatch to monitor the container. The DevOps engineer can use the `AmazonECSContainerDefinition` resource in the CloudFormation template to define the container.

The container definition must include the following properties:

```
"logConfiguration": {
  "logDriver": "awslogs",
  "options": {
    "awslogs-create-group": "true",
    "awslogs-group": "awslogs-mytask",
    "awslogs-region": "us-west-2",
    "awslogs-stream-prefix": "awslogs-mytask",
    "mode": "non-blocking",
    "max-buffer-size": "25m"
  }
}
```


D. Route 53 DNS records are stored in a single Amazon Route 53 zone. DNS records are stored in a single Amazon Route 53 zone. Route 53 records are stored in a single Amazon Route 53 zone. Route 53 records are stored in a single Amazon Route 53 zone. Route 53 records are stored in a single Amazon Route 53 zone.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 140

Amazon Elastic Block Store (Amazon EBS) is used to store data for Amazon EC2 instances. Amazon EBS is used to store data for Amazon EC2 instances. Amazon EBS is used to store data for Amazon EC2 instances. Amazon EBS is used to store data for Amazon EC2 instances.

- A. Amazon EventBridge can be used to trigger AWS Lambda functions when an Amazon EC2 instance is stopped. Amazon EventBridge can be used to trigger AWS Lambda functions when an Amazon EC2 instance is stopped. Amazon EventBridge can be used to trigger AWS Lambda functions when an Amazon EC2 instance is stopped.
- B. Amazon EC2 Auto Recovery can be used to automatically restart Amazon EC2 instances that are stopped. Amazon EC2 Auto Recovery can be used to automatically restart Amazon EC2 instances that are stopped. Amazon EC2 Auto Recovery can be used to automatically restart Amazon EC2 instances that are stopped.
- C. Amazon CloudWatch can be used to monitor the health of Amazon EC2 instances. Amazon CloudWatch can be used to monitor the health of Amazon EC2 instances. Amazon CloudWatch can be used to monitor the health of Amazon EC2 instances.
- D. Amazon Systems Manager Automation Runbooks can be used to automate remediation actions for Amazon EC2 instances. Amazon Systems Manager Automation Runbooks can be used to automate remediation actions for Amazon EC2 instances. Amazon Systems Manager Automation Runbooks can be used to automate remediation actions for Amazon EC2 instances.

Answer: D (LEAVE A REPLY)

<https://aws.amazon.com/blogs/mt/automate-remediation-actions-for-amazon-ec2-notifications-and-beyond- using-ec2-systems-manager-automation-and-aws-health/>

NEW QUESTION: 141

AWS CodeConnections can be used to connect Amazon S3 buckets to Git repositories. AWS CodeConnections can be used to connect Amazon S3 buckets to Git repositories. AWS CodeConnections can be used to connect Amazon S3 buckets to Git repositories.

- A. AWS CodeBuild can be used to build and test code. AWS CodeBuild can be used to build and test code. AWS CodeBuild can be used to build and test code.
- B. AWS CodePipeline can be used to orchestrate the build and deployment of code. AWS CodePipeline can be used to orchestrate the build and deployment of code. AWS CodePipeline can be used to orchestrate the build and deployment of code.
- C. AWS CodeCommit can be used to store and manage source code. AWS CodeCommit can be used to store and manage source code. AWS CodeCommit can be used to store and manage source code.
- D. AWS IAM can be used to manage access to AWS services. AWS IAM can be used to manage access to AWS services. AWS IAM can be used to manage access to AWS services.
- E. AWS CodeBuild can be used to build and test code. AWS CodeBuild can be used to build and test code. AWS CodeBuild can be used to build and test code.
- F. AWS CodeDeploy can be used to deploy code to Amazon EC2 instances. AWS CodeDeploy can be used to deploy code to Amazon EC2 instances. AWS CodeDeploy can be used to deploy code to Amazon EC2 instances.

Answer: C,D,E (LEAVE A REPLY)

NEW QUESTION: 142

Amazon S3 can be used to store data. Amazon S3 can be used to store data. Amazon S3 can be used to store data. Amazon S3 can be used to store data. Amazon S3 can be used to store data.


```
"End": true
}
}
}
```

Create a Development Environment from CloudFormation Template:

Deploy the development environment in a new account using the existing CloudFormation template.

Schedule an EventBridge rule to start the Step Functions state machine on a weekly basis.

EventBridge rule example:

```
{
  "ScheduleExpression": "rate(7 days)",
  "StateMachineArn": "arn:aws:states:<region>:<account-id>:stateMachine:AnonymizeAndCopyData"
}
```

By using Macie for data anonymization and Step Functions for automation, you ensure PII is properly handled before data transfer between environments.

References:

Amazon Macie

AWS Step Functions

AWS CloudFormation Templates

NEW QUESTION: 144

- □□□□ □□□ □□□□ □□□□□(SaaS) □□□□□□□ □□ □□ □□(Proof of Concept)□ □□□□ □□□□. □ □□□□□□□
- AWS Organizations □□ □□□ □□ □□ □□ AWS □□□□ □□□□.
- □□ □□□ □□□ AWS □□□□ □□ □□□ □□ IAM □□□□ □□□□ □□□□. □□□□ □□□□□□ □□□□ □□ □□□ □□□ □□□ □□□ □□□□ □□□□ □□□□.
- □□□□ □□□□ □□□ □□□□ □□□□□?
- A.** □□□□ □□ □□□□ □□□□ IAM □□□□□ □□□□□□. □□□□□ □□□□ □□□□ □□ □□□□ □□ □□ □□□□ □□□□□□. □□ □□□□ □□□□ □□ □□□□ □□□□□□.
- B.** □□□□□ IAM □□□□ □□□□□□. PowerUserAccess □□ □□□□ IAM □□□□ □□□□□□. □□□□ □□□□ □□□□ □□(MFA)□ □□□□□ □□.
- C.** □□□□ □□ □□□□ SCP□ □□□□□□. □□□□□ □□□□ □□□□ □□ iam:*□ □□ □□ □□□□□ SCP□ □□□□□□.
- D.** □□□□□ □□ □□□□ □□□□ □□□□ □□□□ □□□□ IAM □□□□ □□□□ □□ IAM □□□□ □□□□□□. □□□□ □□ □□□□ □□□□□ □□□□□. □□□□□□ □□ □□ □□□□ □□□□□□.

Answer: (SHOW ANSWER)

To allow only creation/configuration of service-linked roles , you need a way to tightly scope what IAM actions the developer can perform.

The most AWS-appropriate mechanism for "delegate limited IAM admin" is a role with a permissions boundary :

* A permissions boundary sets the maximum permissions the role (and any roles it creates, depending on design) can ever have. This is a standard pattern to safely delegate IAM tasks without granting broad IAM administration.

* You can define the role's policy + boundary so the developer can call only the APIs required for service-linked roles (for example actions like creating service-linked roles and passing only allowed AWS service principals), while preventing general role/policy creation outside that scope.

Why the others don't meet "service-linked roles only":

* A is vague ("common services") and cross-account access doesn't inherently restrict to only service-linked roles. It's also more operational overhead than needed for a single dev account use case.

* B PowerUserAccess is far too broad and does not restrict to service-linked roles.

* C An SCP with Deny iam:* would block IAM entirely (including what the developer needs). Even if refined, SCPs set account-wide guardrails and are not the right tool to grant a single developer the precise ability to create only service-linked roles.

NEW QUESTION: 145

□□□□ MySQL □□ Amazon Aurora □□ AZ DB □□□□□□ □□□□□□□□ □□□□ □□□□□□□□ □□□□. □□ □□□ □□ □ □ □ □ □□□□ □□□□□□□□. DevOps □□□□□□ □□ □□ □ □ □ □□□□□□ □□□□□□ □□□ □□□ □□□ □□□□□ □ □□□.

□□ □□□□ □□ □□□□□?

A. □□ □□ □□ Amazon Route 53 CNAME□ □□ □□□ □□ □□□□ □□ □□□□□□ □□ □□□□□□ □□ □□□□□ □□ □ □. AWS CloudTrail□ Amazon RDS □□ □□□ □□ Amazon SNS □□□ □□□□ □□ □□□ □□□ □□□ □□□□□ □□ □□ □□□ □□□□ AWS Lambda □□□ □□□□□.

B. □□ □□□□□□ □□□□□ □□□□□ Aurora □□□ □□ □□□□□□ □□□□□□. □ □□□□□□ □□□□□□ □□□□□□□ □□□□□. AWS Lambda □□□ □□□□ □□□ □□□□□ □□□□ □□ □□□ □□□□□ □□□□□ □□□ □□ □□□□□□ □□□□□ AWS CloudTrail□ □□□□□□.

C. □□□□□□□ AWS CloudFormation □□□□ □□□□ □□□□ □□□□, □□□□ □□□□ □□□ □□□□□□, □□□□□□□□ □□ □□□ □□□□□ □□□□□ AWS Lambda □□□ □□□□□. □□ □□□□ □□□ □ □ Lambda □□□ □□□□ Amazon CloudWatch □□□ □□□□□.

D. AWS Systems Manager Parameter Store□ Aurora □□□□□□ □□□□□□. □□□□□□ □□□ □□□□ AWS Lambda □□□ □ □□□ □□□ □□□□□ □□□□□ AWS Systems Manager Parameter Store□ □□□ □□□□□ URL□ □□□□□□ Amazon EventBridge □□□□ □□□□□. □□□□□□ □□□ □□□ □□ Parameter Store□□ □□□□□□ □□ □□□□□ □□□□□□□ □□□□□.

Answer: D (LEAVE A REPLY)

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: 146

□ □□□ □□ □□□ Amazon S3 □□□ □□□□□ □□□□□□. □ □□□ □□ □□ □□□□ □□□□ □□□ □□ □□ □□□ □□ □ □□ □□□□ □□□.

□□ □□□□ □□ □□ □□ □□□□ □□□ □□ □□□ □□□ □ □□□□?

A. Amazon S3 □□ □□□ □□□ □□□□□□. □□□ □□□ Amazon Aurora □□□□□□□□ □□□□□□. Aurora □□□□□□□□□ SQL □□□ □□□□ □□□ □□□ □□□ □□□□□.

B. Amazon S3 □□ □□□ □□□ □□□□□□. Amazon Athena□ □□□□ □□□ □□□ □□□ □□ □□□□ □□□□□□. Athena □ □□□□ SQL □□□ □□□□□ □□□ □□□ □□□ □□□□□.

C. □□ S3 □□ □□□ □□□□□ □□ AWS Lambda □□□ □□□□□□. □□□ ID, S3 □□ ID, □□ □□ □□□ □□ □□□ □□□ Amazon Aurora □□□□□□□□ □□□□□□ Lambda □□□ □□□□□□. Aurora □□□□□□□□□ SQL □□□ □□□□□ □□□ □□□ □□□□□□□.

D. S3 can generate server access logs that record detailed information about each request, including requester, bucket, key, operation, time, and status. These logs are written as objects to an S3 bucket. To analyze access patterns, the simplest and most serverless approach is to use Amazon Athena directly on those logs without building ingestion pipelines or databases. Option B enables S3 server access logging and then creates an Athena external table over the log bucket. AWS provides standard log formats and even example schemas for S3 access logs. The analytics team can run ad hoc SQL queries to count the number of accesses per object per time period, filter by user, and perform aggregations, all without provisioning compute or managing databases. Option A requires ingesting logs into Aurora, which adds ETL complexity and ongoing database management. Option C requires a Lambda function for every access event plus DB writes, which is more complex and potentially expensive at scale. Option D uses CloudWatch Logs and Managed Flink, which is more suited for streaming analytics and is significantly more complex than necessary for monthly summary reports. Therefore, Option B provides the required analysis with the least development and operational effort.

Answer: (SHOW ANSWER)

Amazon S3 can generate server access logs that record detailed information about each request, including requester, bucket, key, operation, time, and status. These logs are written as objects to an S3 bucket. To analyze access patterns, the simplest and most serverless approach is to use Amazon Athena directly on those logs without building ingestion pipelines or databases.

Option B enables S3 server access logging and then creates an Athena external table over the log bucket.

AWS provides standard log formats and even example schemas for S3 access logs. The analytics team can run ad hoc SQL queries to count the number of accesses per object per time period, filter by user, and perform aggregations, all without provisioning compute or managing databases.

Option A requires ingesting logs into Aurora, which adds ETL complexity and ongoing database management. Option C requires a Lambda function for every access event plus DB writes, which is more complex and potentially expensive at scale. Option D uses CloudWatch Logs and Managed Flink, which is more suited for streaming analytics and is significantly more complex than necessary for monthly summary reports.

Therefore, Option B provides the required analysis with the least development and operational effort.

NEW QUESTION: 147

AWS Organizations can be used to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it. Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

Options A and B are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account. Option C requires a Lambda function for every access event plus DB writes, which is more complex and potentially expensive at scale. Option D uses CloudWatch Logs and Managed Flink, which is more suited for streaming analytics and is significantly more complex than necessary for monthly summary reports.

A. AWS Organizations can be used to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it. Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

B. AWS Organizations can be used to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it. Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

C. AWS Organizations can be used to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it. Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

D. AWS Organizations can be used to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it. Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

Answer: A (LEAVE A REPLY)

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html?](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html?icmpid=docs_orgs_console)

icmpid=docs_orgs_console

SCPs (Service Control Policies) are the best way to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it.

Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

NEW QUESTION: 148

DevOps □□□□□ □□□□ □□□□ □□□□, □□ ECR □□□□, □□□□ □□□□ □□□□, □□□□ □□ □□□ □□□□ □ □□ □□□□ □□ □□□□ CI/CD □□□□□□ □□□□□□.

□□ □□□□ □□ □□□□□□?

- A. □□ ECR □□□□ □□□, □□ □ □□□□ □□□□, □□ □□□ □□ □□□□ □□□□□ □□□□ □□□□□.
- B. □□□□ □□□□ □□□□ □□□□ □□ □□ ECR □□□□ □□□, □□□□ □□□□ □□ □□□ □□□□ □□□, □□ □ □□ □□ □□□□□.
- C. □□ ECR □□□□ □□□, □□□ □□ □□□ □□□□, □□□□ □□□ □□ □□□□ □□□, □□ □□□□ □□□□□□.
- D. □□ ECR □□□□ □□□, □□ □□□□ □□□□□, □□□ □□ □□□ □□□□.

Answer: D (LEAVE A REPLY)

- * ECR pull-through cache caches images from upstream repositories for resilience.
- * Private repo with basic scanning ensures vulnerability detection on pushed images.
- * Enabling pull-through caching on a private repo combines caching and vulnerability scanning seamlessly.
- * Public repos do not support pull-through caching of upstream images.
- * Replication rules are for multi-Region replication, not upstream caching.

References:

Amazon ECR Pull Through Cache

Amazon ECR Image Scanning

NEW QUESTION: 149

□ □□□ AWS Organizations □ □□□ □□□□ AWS □□□ □□□□□. □□□ □□□ □□□□ □□□ AWS □□□ □□□□ □□□ □ □□□□ □□□□□□ □□□□ □□□□.

□□□□ □□ □ □□□ □□□□ □□□□ □□ □□□ □□ □□□□. □□□ □□ □□□ □□□ □□□ □□□□ □□□□□. □□□ □□ □□□□□□ □□□ □ □□□ □□□ □ □□□ □□□□ AWS CloudTrail □□□□ □□□□□ □□□.

□□□□□ □ □□□□ □□□□ □□□ □□□ □□□ □□□□ □□□?

- A. □□□ □□□ Amazon EventBridge □□□ □□□□ □□ □□ □□□□ □□□ □□□ □□ □□□ □□□ □□□□□. □□□ □□□ □□□□ □□□□□ □□□ □□□ □□ □□□ □□□ □□□□□□□.
- B. □□□ □□□ □□□ □□ Amazon EventBridge □□□ □□□ □□□□□. □□□ □□□ □□□□ □□□□□ □□□ □□ □□□ □ □□ □□□□□□□. □□□ □□□ CloudTrail □□□□ □□ □□□□□ EventBridge □□□ □□□ □□□ □□□□□.
- C. □□□ □□□ □□□ □□□□ □□□ □□ Amazon EventBridge □□□ □□□ □□□□.
- □□□ □□□ □□□ □□ □□□ □□ □□□ □□□ □□□□□ EventBridge □□□ □□□ □□□□.
- D. □□□ □□□ □□□ □□ Amazon EventBridge □□□ □□□ □□□□□. □□□ □□ □□□ □□□ □□□ □□□ □□ □□□ □ □□ □□□□□ EventBridge □□□ □□□ □□□□□.

Answer: A (LEAVE A REPLY)

Use cross-account EventBridge by configuring a rule in the source (automation) account to send events to the target accounts' default event buses, and grant permissions on the target default event buses to accept events from the source account. This is the standard cross-account event routing model.

NEW QUESTION: 150

□ □□□ Amazon EKS □□□□□ □□□□ □□□ □□ □□□□ □□□ □□ □□□□ □□□ □□□□ □□□. □□, API □□□ □□ □□ □□□□ □□□ □□□□□□ □□□.

□□ □□□□ □□ □□□ □□□□□□ □□□ □□ □□□ □□□ □ □□□□?

- A. AWS CloudTrail, Logstash.
- B. CloudWatch, Pod, CloudWatch Container Insights.
- C. S3 API, Kubernetes.
- D. OpenTelemetry, AWS Distro, Amazon Redshift.

Answer: B (LEAVE A REPLY)

Enabling EKS control plane logs to CloudWatch captures API requests. CloudWatch Container Insights collects node and pod-level performance data with no additional infrastructure. AWS recommends this integrated observability solution for minimal management overhead.

NEW QUESTION: 151

SaaS ALB, CodePipeline + CodeDeploy, ECS(Fargate). Traffic increments by linearPercentage every linearInterval until 100%. This provides zero-downtime gradual rollout - as per CodeDeploy ECS Blue/Green Traffic Shifting documentation.

- A. appspec.yaml, TimeBasedLinear.
- B. AllAtOnce.
- C. TimeBasedCanary.
- D. ALB, ECS(Fargate).

Answer: (SHOW ANSWER)

CodeDeploy supports TimeBasedLinear traffic shifting for ECS blue/green deployments. Traffic increments by linearPercentage every linearInterval until 100%. This provides zero-downtime gradual rollout - as per CodeDeploy ECS Blue/Green Traffic Shifting documentation.

DOP-C02-KR DumpTop DOP-C02-KR! DumpTop DOP-C02-KR, DumpTop DOP-C02-KR. <https://www.dumptop.com/Amazon/DOP-C02-KR-dump.html> (439 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)

NEW QUESTION: 152

- A. Amazon DynamoDB.
- B. Aurora, Amazon DynamoDB.
- C. Aurora, Amazon Redshift.
- D. Amazon Redshift, Amazon DynamoDB.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 153

DevOps wants to build a pipeline that builds and deploys an application to Amazon S3. The pipeline should build the application and upload the artifacts to Amazon S3. The pipeline should be able to build the application and upload the artifacts to Amazon S3. The pipeline should be able to build the application and upload the artifacts to Amazon S3.

- A. AWS CodeArtifact. CodeArtifact is a managed artifact repository that stores and distributes artifacts. It is used to store and distribute artifacts. It is used to store and distribute artifacts.
- B. Amazon Elastic Container Registry(Amazon ECR). Amazon ECR is a managed container registry that stores and distributes container images. It is used to store and distribute container images. It is used to store and distribute container images.
- C. EC2 Image Builder. EC2 Image Builder is a managed service that builds and distributes Amazon EC2 images. It is used to build and distribute Amazon EC2 images. It is used to build and distribute Amazon EC2 images.
- D. Amazon S3. Amazon S3 is a managed storage service that stores and distributes objects. It is used to store and distribute objects. It is used to store and distribute objects.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 154

A DevOps engineer wants to build a pipeline that builds and deploys an application to Amazon S3. The pipeline should build the application and upload the artifacts to Amazon S3. The pipeline should be able to build the application and upload the artifacts to Amazon S3. The pipeline should be able to build the application and upload the artifacts to Amazon S3.

```
version: 0.2
phases:
  build:
    commands:
      - go build -o myapp
  post_build:
    commands:
      - aws s3 cp --acl authenticated-read myapp s3://artifacts/
```

- A. post_build commands: aws s3 cp --acl public-read myapp s3://artifacts/
- B. build commands: aws s3 cp --acl public-read myapp s3://artifacts/
- C. build commands: aws s3 cp --acl authenticated-read myapp s3://artifacts/
- D. post_build commands: aws s3 cp --acl authenticated-read myapp s3://artifacts/

Answer: D (LEAVE A REPLY)

When setting the flag authenticated-read in the command line, the owner gets FULL_CONTROL. The AuthenticatedUsers group (Anyone with an AWS account) gets READ access. Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/acl-overview.html>

NEW QUESTION: 155

- Which AWS service is best suited for running mobile app tests across a wide range of devices and operating systems?
- A. AWS Device Farm
 - B. Amazon EC2
 - C. Amazon Elastic Container Service (Amazon ECS)
 - D. AWS Lambda

Answer: (SHOW ANSWER)

AWS Device Farm provides a fully managed environment for automated and manual app testing across real physical devices. It integrates directly with AWS CodePipeline and automatically scales based on test concurrency. AWS documentation recommends Device Farm for mobile app testing due to its scalability and zero infrastructure management.

NEW QUESTION: 156

- Which AWS service is best suited for analyzing and finding the number of logins for a specific user from the past 7 days?
- A. Amazon CloudWatch Logs
 - B. Amazon CloudWatch Logs Insights
 - C. Amazon CloudWatch Logs Insights
 - D. Amazon CloudWatch

Answer: C (LEAVE A REPLY)

To analyze and find the number of logins for a specific user from the past 7 days, a CloudWatch Logs Insights query is the most suitable solution. CloudWatch Logs Insights enables you to interactively search and analyze your log data in Amazon CloudWatch Logs. You can use the query language to perform queries that contain multiple commands, including aggregation functions, which can count the occurrences of logins for a specific username over a specified time period. This approach is more direct and efficient than creating a metric filter or subscription, which would require additional steps to publish and sum a metric. References: AWS Certified DevOps Engineer - Professional, CloudWatch Logs Insights query syntax, Tutorial: Run a query with an aggregation function, Add or remove a number widget from a CloudWatch dashboard.

NEW QUESTION: 157

- Which AWS service is best suited for storing and retrieving Docker images?
- A. Amazon Elastic Container Registry (Amazon ECR)
 - B. Amazon Elastic Container Service (Amazon ECS)
 - C. Amazon CodeBuild
 - D. Amazon CodeDeploy

DevOps □□□□□ CI/CD □□□□□□ □□□ □□□ □□□□ □□□. CI/CD □□□□□□ CRITICAL □ HIGH □□□ □□ □□□□ □□□□ □□□ □□□□ □□□.

□□□ □□ □□□ □□□□□ □□ □□ □□□ □□□□□? (□ □□□ □□□□□.)

A. Amazon ECR □□ □□□□ □□□□□□.

B. Amazon ECR □□□ □□□□ □□□□□□.

C. □□□ □□□□ CRITICAL □□ HIGH □□□ □□□□ CI/CD □□□□□□ □□□ □□□ □□□□□ Amazon ECR□ □□□□□.

D. □□□ □□□ □□□□ AWS Lambda □□□ □□□□□ Amazon EventBridge □□□ □□□□□. Lambda □□□ Amazon Inspector □□ □□□ □□□□□ CI/CD □□□□□□ □□ □□ □□ □□□ □□□□□ □□□□□.

E. □□□ □□□ □□□□ AWS Lambda □□□ □□□□□ Amazon EventBridge □□□ □□□□□. Lambda □□□ Clair □□ □□□ □□□□□ CI/CD □□□□□□ □□ □□ □□ □□□ □□□□□ □□□□□.

Answer: B,D (LEAVE A REPLY)

Amazon ECR supports enhanced scanning powered by Amazon Inspector, which provides deeper security scanning for container images including OS and programming language package vulnerabilities.

* Enabling enhanced scanning (Option B) allows detection of CRITICAL and HIGH vulnerabilities.

* Amazon ECR emits scan completion events via EventBridge, which can trigger Lambda functions. The Lambda function can process the scan results from Amazon Inspector and programmatically approve or reject the image in the CI/CD pipeline (Option D).

* Basic scanning (Option A) is limited and does not integrate with Inspector.

* Options C and E describe functionalities not natively supported (ECR does not automatically submit Rejected status; Clair is not used in AWS ECR scanning).

Reference:

Amazon ECR Enhanced Scanning: " Enhanced scanning powered by Amazon Inspector identifies vulnerabilities in container images.

" (Amazon ECR Image Scanning) Using EventBridge and Lambda for Scan Status: " ECR emits scan events that can be used to trigger Lambda functions for custom approval workflows. " (Amazon ECR Scan EventBridge)

NEW QUESTION: 158

DevOps □□□□□ Amazon EC2□□ □□□□ □□□ □□ □ □□□□ □□□□□. □ □□□□ Amazon Kinesis Data Streams□ □□ □□ □ □□□ □□□□ □□□□□. DevOps □□□□□ Amazon EC2□□□ □□□□□ Kinesis □□□ □□□□□□□ □□□□□.

□□□□ □□□ □□□□ Kinesis □□□ □□□□□□□ □□□□ □□□□ □□ □□ □□□□ Kinesis □□□ □□□□ □□□□ □□ □□□.

DevOps □□□□□ □□□ □□□ □□□□ □□ □□□□ □□□□ □□□.

□□ □□ □□□□ □□ □□□ □□ □□□ □□□□ □□□□ □□□□□? □□□ Amazon S3□ □□□ □□□□□ □□□□□ Kinesis □□□ □□□□□□□ □□□□□. Amazon EMR□ □□□□ Amazon S3□□ □□ □□□□ □□□□ □□ □□□□ □□□□. □□□ Amazon S3□ □□□□□.

B. Amazon CloudWatch GetRecords IteratorAgeMilliseconds □□□ □□□□ □ □□ EC2 □□□□□ □□□□□ Kinesis □□□ □□□□ □□□ □□□□ □□□□□. Kinesis □□□ □□□□ □□ □□□ □□□□.

C. AWS Lambda □□□ □□□□□ Kinesis □□□ □□□□□□□ □□□□□. □□□ □□□□ □□□□ □□ Lambda □□□ □□□ □□□ Kinesis □□□ □□□□ □□□□□.

D. □□□ □□□□□□□ □□□□ □ □□□ □□□ □ □□□ Kinesis □□□ □□□□ □□ □□ □□ □□ □□□□ □□□□.

Answer: B (LEAVE A REPLY)

<https://docs.aws.amazon.com/streams/latest/dev/monitoring-with-cloudwatch.html> GetRecords.IteratorAgeMilliseconds - The age of the last record in all GetRecords calls made against a Kinesis stream, measured over the specified time period. Age is the difference between the current time and when the last record of the GetRecords call was written to the stream. The Minimum and Maximum statistics can be

used to track the progress of Kinesis consumer applications. A value of zero indicates that the records being read are completely caught up.

NEW QUESTION: 159

Which AWS service can be used to centrally manage AWS IoT Greengrass devices? (3 correct answers.)

Options:

- A. Amazon EC2
- B. AWS IoT Greengrass
- C. AWS Systems Manager
- D. Amazon EventBridge
- E. AWS IAM
- F. Amazon SSM

Answer: C,E,F (LEAVE A REPLY)

https://aws.amazon.com/blogs/mt/how-to-centrally-manage-aws-iot-greengrass-devices-using-aws-systems-manager/?force_isolation=true

NEW QUESTION: 160

Which AWS service can be used to centrally manage AWS CodePipeline pipelines? (2 correct answers.)

Options:

- A. Amazon CloudWatch Logs
- B. Amazon S3
- C. AWS CloudTrail
- D. AWS Lambda
- E. Amazon SSM

Answer: (SHOW ANSWER)

To meet the new guideline for application deployment, the company can use a combination of AWS CodePipeline and AWS CloudTrail. A manual approval action in CodePipeline allows the security team to review and approve changes before they are deployed. This action can be configured to pause the pipeline until approval is granted, ensuring that no changes move to production without the necessary sign-off.

Additionally, by creating an AWS CloudTrail trail, all actions taken within CodePipeline, including approvals, are recorded and delivered to an Amazon S3 bucket. This provides an audit trail that can be retained for compliance and review purposes.

AWS CodePipeline's manual approval action provides a way to ensure that a member of the security team can review and approve changes before they are deployed¹.

AWS CloudTrail integration with CodePipeline allows for the recording and retention of all pipeline actions, including approvals, which can be stored in Amazon S3 for record-keeping².

NEW QUESTION: 161

A company is migrating its application development workflow to AWS CodeBuild. The application code is currently stored in a Git repository. The company wants to ensure that the CodeBuild project can clone the repository and run the tests. Which of the following configurations is the most appropriate?

A. CodeBuild project is configured to use the AWS CLI to clone the repository and run the tests. The CodeBuild project is also configured to use the AWS CLI to push the test results to an Amazon S3 bucket.

B. CodeBuild project is configured to use native Git to clone the repository and run the tests.

C. CodeBuild project is configured to use native Git to clone the repository and run the tests. The CodeBuild project is also configured to use the AWS CLI to push the test results to an Amazon S3 bucket.

D. CodeBuild project is configured to use the AWS CLI to clone the repository and run the tests. The CodeBuild project is also configured to use the AWS CLI to push the test results to an Amazon S3 bucket.

E. CodeBuild project is configured to use native Git to clone the repository and run the tests. The CodeBuild project is also configured to use the AWS CLI to push the test results to an Amazon S3 bucket.

F. CodeBuild project is configured to use native Git to clone the repository and run the tests. The CodeBuild project is also configured to use the AWS CLI to push the test results to an Amazon S3 bucket.

Answer: (SHOW ANSWER)

Step 1: Using Native Git in CodeBuild To meet the requirement of running unit tests and tagging the most recent commit if the tests pass, the CodeBuild project should be configured to use native Git to clone the CodeCommit repository. Native Git support allows full functionality for managing the repository, including the ability to create and push tags.

Action: Configure the CodeBuild project to use native Git to clone the repository and run the tests.

Why: Using native Git provides flexibility for managing tags and other repository operations after the tests are successfully executed.

Step 2: Tagging the Most Recent Commit Once the unit tests pass, the CodeBuild project can use native Git to create a tag for the most recent commit and push that tag to the repository. This ensures that the tagged commit is linked to the test results.

Action: Configure the project to use native Git to create and push a tag to the repository if the tests pass.

Why: This ensures the correct commit is tagged automatically, streamlining the workflow.

Reference: AWS documentation on AWS CodeBuild and Git integration.

This corresponds to Option A: Configure the CodeBuild project to use native Git to clone the CodeCommit repository. Configure the project to run the unit tests. Configure the project to use native Git to create a tag and to push the Git tag to the repository if the code passes the unit tests.

NEW QUESTION: 162

□ □□□ □□□□ □□□□ □□□□ □□ □□ □□ □□□□□□□□ □□□□□□ □□□□. □ □□□□□□□□□ □□ □□ □□□□□□ □□ □□□□, □□ □□ Amazon RDS DB □□□□□□ □□□□ □□ □□□ □□ □□□ □□□□□□. DevOps □□□□□□ □□□□□□ □□□□□□ □□□□ □□□□.

□□ □□□□ □ □□ □□□ □□□□□□?

- A. RDS DB □□□□□□ □□ □□□□□□ □□□□ □□ □□□ □□□ □□ □□□□ □□□□□□.
- B. □□□□□□□□ Amazon DynamoDB □□□□ □□□□ □□□□□□□□□□. Amazon Route 53 □□ □□□ □□□□□ AWS □□ □ □□ □□ □□□ □□□□□□.
- C. RDS DB □□□□□□ □□ AZ □□□ □□□□□□. □□ □□□□□□ □□□ □ □□ □□ □□ □□□□□□ □□ □□ □□(failover)□ □□ □□□ □□□□□□.
- D. □□□□□□□□ □□ □□ □□□ Amazon EC2 □□□□□□ □□□□□□□□□□□□. Amazon Elastic Block Store(Amazon EBS) □□ □□ □ □□□□ □□ □□□□□□ □□ EBS □□□ □□□□□□.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 163

DevOps □□□□□□ □□□□□□□□ □□□□ □□□□ □□□□. □□□□□□□□ Amazon EC2 □□□□□□ □□□□ Amazon Elastic Kubernetes Service(Amazon EKS) □□□□□□ □□□□□□ □□□□. EC2 □□□□□□ Amazon Elastic File System(Amazon EFS) □□ □ □□□ □□□□ □□□□ □□□□ □□□□. Amazon EFS □□□□□ □□□□□ □□□□□□(CSI) □□□□□□ EKS □□□□□□ □□□□□ □□ □□.

DevOps □□□□□□ □□□□□□□□□□ □□□□□□ EC2 □□□□□□ EFS □□ □□□□ □□□□□□ □□□□□□.

□□ □□□□ □□□ □□□ □ □□□□□? (□ □□□ □□□□□□.)

- A. EKS □□□ Amazon EC2□□□ AWS Fargate□ □□□□□□.
- B. EKS □□□□□□□□ NFS □□□□□ □□□□ □□ EFS □□ □□□□ □□ □□□ □□□□ □□□ □□□□□□.
- C. Amazon EFS CSI □□□□□□ □□ □□□□□ □□ □□□ □ □□□ □□ IAM □□□ □□□□□□.
- D. EFS □□ □□□□□□ EKS □□ □ □□ □□□ □□□□□ □□ AWS DataSync□ □□□□□□.
- E. EKS □□□ □□□□□ □□ EFS □□ □□□□□ □□ □□□ □□□ □□□□□.
- F. EFS □□ □□□□ □□□□□ □□□□□□□□□□.

Answer: (SHOW ANSWER)

Mounting EFS to EC2-backed EKS nodes requires:

- * NFS (port 2049) open from nodes to EFS (Security Group rule).
- * Mount targets in each subnet/AZ where nodes reside.
- * IAM role for the EFS CSI driver with elasticfilesystem:ClientMount and ClientRootAccess permissions. These are the standard setup requirements in "Using the Amazon EFS CSI Driver with Amazon EKS."

NEW QUESTION: 164

□□□□ AWS □□ □□□ □□□□ □□ bash □□ □□□□□ □□□□□ □□□. □ □□□ □□ ALB(Application Load Balancer) □□ Amazon EC2 □□□□ □□□ LAMP □□□□□□□ □□□□ □□□□. □□□□ □□ □□ □□□ □□□ □□□□□□□ □□□□□, □□□□ □□ □ □□□□, □□ □□□□ □□□□□ □□ □□ □ □□□□□□, □□ □□□ □□□□□□□□. □□□□ AWS □□□ □□□ □□ □□□ □□ □□□ □□ □□□ □□□□□ □□□.

□□□ □□ □□□ □□□□ □□□□ □□□□□?

- A. AWS CodeBuild □□□□ □□□□□□□ □□□□□□□. AWS CodeDeploy □ appspec.yml □□□□ □□□ bash □□□□□ □□ □□ □□□□ □□ □□□□ ALB □□□□□□ □□ □□ □ □□□□□□. □□□□ □□ □□□□ □□ □□ □□□ □□□□□□□□ appspec.yml □□□□ □□□□□□□.
- B. AWS CodePipeline □ □□□□ AWS CodeCommit □□□□□□□ AWS CodeDeploy □ □□□□□□□ □□□□□□. CodeDeploy □ □ □□□ □□□□ □□□□□□□□ □□□□□ □□□□□□□□ ALB □□ □□ □ □□□□□□□. □□□□ □□ □□□□□□. □□□□ □□ □ □□□ □□ □□ □□□ □□□□□□□□ appspec.yml □□□□ □□□□□□□.
- C. AWS CodePipeline □ □□□□ AWS CodeCommit □□□□□□□ AWS CodeDeploy □ □□□□□□□ □□ □□□ □□□□□□. CodeDeploy □ □□□□ □□□□□□□□ □□□□□□□. CodeDeploy □ appspec.yml □□□ □□□□□ □□□ □□ □□□□ □□ □□□□□ □□ □□□□□ □□□ □□□□□□□□. AWS CodeBuild □ □□□□□ ALB □□□□□□ □□ □□□□□ □□ □□□□□□□.
- D. AWS CodePipeline □ □□□□ AWS CodeBuild □ □□□□□□ □□□□□□□□ □□□□□□□. AWS CodeDeploy □ appspec.yml □□□ □ □□□ bash □□□□□□ □□□□ □□□□ □□ □□□□□□. ALB □□□□□ AWS CodeDeploy □□ □□□□ □□□□□□ □□ □□□ □ □□ □□□□□□. □□□□ □□ □□□□□ □□ □□ □□□ □□□□□□□□ appspec.yml □□□□ □□□□□□□□□□.

Answer: (SHOW ANSWER)

<https://aws.amazon.com/blogs/devops/how-to-test-and-debug-aws-codedeploy-locally-before-you-ship-your-code/#:~:text=You%20can%20test%20application%20code,local%20server%20or%20EC2%20instance.>

NEW QUESTION: 165

□□□ □□□ □□ Amazon Cloud Formation □□□□ □□□□ □□ □□ □□□□□□□□ □□□□□□. □□□□ □□□□ □□ □□□ □ □□□ □□□. □□□□ □□□ □□□□ □□□ □ □□ □□□□ □□□□ □□□ □ □□□□ □□ □□□□□ □□□□□ □□□□□. □□ □□□□ □□□□□ □□ □□□□ □□ □□ □□ □□□ □□□ □□□.

□□□ □□□ □□□□ □□ □□□ □□□ □□ □□□?

- A. □□□ □□□ Cloud Formation □□□□ □□□□ AWS Lambda □□□ □□□□□□. □□ □□□ □□□□ □□□ □□□□□ □□ □□□.
- B. Amazon S3 □□ Cloud Formation □□□ □□□ Amazon S3 □□□□ □□□□ CloudFormation □□□□ □ Amazon SNS □□□ □ □□□□□□□.
- C. CloudFormation StackSets □ □□□□ □□□□ □□□ □□□□ □□□ □□□□□ □□ □□ □□□□ □□□ □□□□□□□.
- D. CloudFormation □□ □□ □ □□ □□ □□(□□ □□ Amazon SNS □ □□□□ □□□ □□□□□ □□ □□□□□.

Answer: (SHOW ANSWER)

This solution will meet the requirements because it will use CloudFormation nested stacks and stack sets to deploy the templates more efficiently and consistently across multiple regions. Nested stacks allow the company to separate out common components and reuse templates, while stack sets allow the company to create stacks in multiple accounts and regions with a single template. The company can also use Amazon SNS to send notifications to the data engineering team whenever a change is made to the templates or the stacks. Amazon SNS is a service that allows you to publish messages to subscribers, such as email addresses, phone numbers, or other AWS services. By using Amazon SNS, the company can ensure that the data engineering team is aware of all changes to the templates and can take appropriate actions if needed. What is Amazon SNS? - Amazon Simple Notification Service

□□□ □□ □□□ EC2 □□□□□ □□ □□□□□□ □□ □□ EC2 □□□□□ □□ □□□ □□□□□ □□ □□□□ □ □□ □□ □ □□□ □□□. DevOps □□□□□ □□□ □□ □□□ □□□□ □□ □□□ □□ □□□?

- A. □□ □□□ Auto Scaling □□□ □□ □□ □□□□ □□□□ □□ □□□ EC2 □□□□□ □ □□□□□□ □□□□□. □□ □□□ □ □□□□ AWS CLI □□□ □□□□ ALB□ □□□□□□ □□□□ □□□□ □□□□. □□□ □□ □□.
- B. AWS CLI □□□ □□□□ ALB□ □□□□□□ □□ □□□ □□ □□□□ □□□□ □□□□. □□ □□ □□ □□□ □□ Auto Scaling □□□ □□ □□□□ □□□□ □□ □□□ EC2 □□□□□ □ □□□□□□ □□□□□.
- C. □□ □□□□ □□□□□□ □□ □□□ EC2 □□□□□ □□ □□□ □□□□□□ □□□□□. □ □□□□ □□ □□□ Auto Scaling □□□ □□□□ □□ □□□□□. □□ □□□ EC2 □□□□□ □□□□□ □□ □□□□□.
- D. □□ □□□ □□ Auto Scaling □□□ □□ □□□□ □□□□ □□ □□□ EC2 □□□□□ □ □□□□□□ □□ □□ □□□□ □□ □□ ALB□ □□ □□ □□□□□□ □□□□□ Route 53 DNS□ □□□□□□□. .

Answer: (SHOW ANSWER)

This solution will meet the requirements because it will use a rolling restart to gradually replace the EC2 instances in the green environment with new instances that have the new software version installed. A rolling restart is a process that terminates and launches instances in batches, ensuring that there is always a minimum number of healthy instances in service. This way, the green environment can be updated without affecting the availability or performance of the application. When the rolling restart is complete, the DevOps engineer can use an AWS CLI command to modify the listener rules of the ALB and change the default action to forward traffic to the green environment's target group. This will switch the traffic from the blue environment to the green environment all at once, as required by the question.

NEW QUESTION: 168

□□□ □□□□ □□□□ AWS □□ □□□ □□ □□□ □□□ □□ □□□□□ □□□. □□ □□□ AWS Management Console□□ □ □□□ □□□□□. □□□□ □□ AWS CloudFormation□ □□□□ □□□□ □□□□ □□□□ Amazon VPC □ □□ □□□□ □□ □ □ □□ □□ □□□□□. □□ □□□□ Application Load Balancer, Amazon EC2 Auto Scaling □□, □□ □□ □ Amazon DynamoDB □ □□□ □□ □□ □□ □□□□.

□□□ □□□□ □□ DevOps □□□□□ □□ □□ □□□ □□□□□□ □□□. □□□□□□□ □□□□ □ □□□ □□□□ □□□ □□□ □□□□ □□□ □□□ □□□ □□ □□□□□ □ □□ □□□ □□□ □□□. CloudFormation□ □□ □□□ □□□□ □□□ □ □ □□□□□.

□□□ □□ □□□ □□□□ □□□□□ □□□ AWS □□□ □□□□ □□□□ □□ □□□ □□□□□?

- A. □□□□ □□□ □□□□ Fn::ImportValue □□ □□□ □□□□ Virtual Private Cloud(VPC) □ □□□ □□ □□□□□. Count □□ □ □ □□□ □□□□ □□□ □□ □□ □□□□ CloudFormation StackSets□ □□ □□□ □□□□□□□. UpdateStackSet □□□ □□□□ □□ □□ □□□ □□□□□□□.
- B. □□ □□□ □□□□ □□ □□□ □□ □□□ □□□□□. □□□ □□ □□□□□□□ TemplateURL□ □□□□ □□□□ □□ □□□ □ □□□□□□. Virtual Private Cloud(VPC) □ □□□ □□ □□□□□ □□ □□□□ □□□□ □□□□ Fn::ImportValue □□ □□□ □ □□□□□. CreateChangeSet □ ExecuteChangeSet □□□ □□□□ □□ □□ □□□ □□□□□□□.
- C. □□ □□□ □□□□ □□ □□□ □□ □□□ □□□□□. Fn::ImportValue □□ □□□ □□ □□□ □□□□ □□ □□□□ Virtual Private Cloud(VPC) □ □□□ □□ □□□□□. CreateChangeSet □ ExecuteChangeSet □□□ □□□□ □□ □□ □□□ □□□□□□ □.
- D. □□ □□□□ □□□□ □□□□ Fn::ImportValue □□ □□□ □□□□ Virtual Private Cloud(VPC) □ □□□ □□ □□□□□. CloudFormation □□ □□□□ □□□□ □□ □□□□ □□ □□□□ □□□□□. CreateChangeSet□ □□□□□. □ ExecuteChangeSet □□□ □□□□ □□ □□ □□□ □□□□□□□.

Answer: C (LEAVE A REPLY)

<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: 169

□□□ AWS Lambda □□□ □□□ □□□□□□□ □□□□. Lambda □□□ AWS CodeCommit □□□□□□ □□□ Python □□□ □ □□□□. □ □□□ □□ Python □□ □□□ □□ □□□□ □□□□ □□□ □□□□□□. □□□□□ □□□□ □□□ □□ □□□ □ □□□ □□ □□□□ □□ Lambda □□□ □□ □□ □□□□ □□□□□□.

□□□ DevOps □□ □□ □□□□ □□ AWS CodePipeline □□□□□□ □□□□ □□ □□□□ □□□□ □□□. □□□□ □□□ □ □□ □□ □□□□ □□ □□□□ □□□□ □□□.

□□ □□□□ □□□ □□ □□□ □□□□□□? **A.** CodeCommit □□□□□□ Amazon CodeGuru □□□□ □□□□□□. □□□ AWS CodeBuild □□□□□ □□□□□□. CodePipeline □ □□□□□□ □ CodeBuild □□□□□ □□□□ □□□ □□□ □□□□□□. CodeCommit □□□□□□ buildspec.yml □□□ □□□□ □. buildspec.yml □□□□ CodeGuru □□□ □□□□ □□□ □□□□□□.

B. □ AWS CodeBuild □□□□□ □□□□□□. CodePipeline □□□□□□□ □ CodeBuild □□□□□ □□□□ □□□ □□□ □□□□ □. CodeBuild □□□ □□□ □□□□□□. CodeCommit □□□□□□ buildspec.yml □□□ □□□□□□. buildspec.yml □□□□ □□ □□ □□□ JUNITXML □□□□ □□ □□□□ □□□□ □□□ □□□□□□. □ CodeBuild □□□ □□□ □□□□ □□□ □□□□ □□□□ □.

C. □□□ AWS CodeArtifact □□□□□□ □□□□□□. □□□ AWS CodeBuild □□□□□ □□□□□□. CodePipeline □□□□□□□ □ CodeBuild □□□□□ □□□□ □□□ □□□ □□□□□□. □□ CodeCommit □□□□□□ appspec.yml □□□ □□□□□□. appspec.yml □□□ □□ □□ □□□□ CUCUMBERJSON □□□□ □□ □□□□ □□□□ □□□ □□□□□□. □ CodeArtifact □□□ □□□ □□□□□ □□□ □□□□□ □□□□□□.

D. □ AWS CodeBuild □□□□□ □□□□□□. CodePipeline □□□□□□□ □ CodeBuild □□□□□ □□□□ □□□ □□□ □□□□ □. □ Amazon S3 □□□ □□□□□□. CodeCommit □□□□□□ buildspec.yml □□□ □□□□□□. buildspec.yml □□□ □□ □□□□ HTML □□□□ □□ □□□□ □□□□ □□□ □□□□□□. □□□ □□□□ □□□ □□□□□ S3 □□□ □□□□□□.

Answer: B (LEAVE A REPLY)

The correct answer is B. Creating a new AWS CodeBuild project and configuring a test stage in the AWS CodePipeline pipeline that uses the new CodeBuild project is the best way to integrate the unit tests into the existing pipeline. Creating a CodeBuild report group and uploading the test reports to the new CodeBuild report group will produce reports about the unit tests for the company to view. Using JUNITXML as the output format for the unit tests is supported by CodeBuild and will generate a valid report.

Option A is incorrect because Amazon CodeGuru Reviewer is a service that provides automated code reviews and recommendations for improving code quality and performance. It is not a tool for running unit tests or producing test reports. Therefore, option A will not meet the requirements.

Option C is incorrect because AWS CodeArtifact is a service that provides secure, scalable, and cost-effective artifact management for software development. It is not a tool for running unit tests or producing test reports.

Moreover, option C uses CUCUMBERJSON as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.

Option D is incorrect because uploading the test reports to an Amazon S3 bucket is not the best way to produce reports about the unit tests for the company to view. CodeBuild has a built-in feature to create and manage test reports, which is more convenient and efficient

than using S3. Furthermore, option D uses HTML as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.

NEW QUESTION: 170

- Windows □ Linux Amazon EC2 □□□□□□ □□□□□□ □□□□□. □□□□□ AWS □□□ □□ □□ □□□□ □□□□ □. □□□ □ □□□□□□□ □□ Auto Scaling □□□ □□□□□.
- □□□□□ □□ □□□ □□ □□□□ □□□□□ □□□□□. □□□□□ Windows□ SMB□ □□□□ □□ Linux□ NFS□ □□□ □ □□□. □□ □□□□ □□ □□□ □□□ □□□□□ □□□. □□ □□□□□ □□□□ □□ □□□.
- □□ □□□ □□□□ □□ □□□ □□□□□? (3□□ □□□□□.)
- A. □□ □□ □□□ □□□ □□ Amazon Elastic File System(Amazon EFS) □□ □□□□ □□□□□.
 - B. NetApp ONTAP □□ AZ □□ □□□□ Amazon FSx□ □□□□□.
 - C. □□ □□□□□ □□□ □□ SSD(gp3) Amazon Elastic Block Store(Amazon EBS) □□□ □□□□□.
 - D. □□ □□□□ □□□□ □□ □ □□□□□□□□ □□ □□□□ □□ □□□ □□□□□□□□□.
 - E. □ Auto Scaling □□□□ □□□□ □□ □□□ □□□□□.
 - F. □ □□□□□ □□□ □ □□ □□□□ □□□□□ □ □□□□□□□□□ EC2 □□□□□ □□□□□□□□□.

Answer: A,B,D (LEAVE A REPLY)

Create an Amazon Elastic File System (Amazon EFS) File System with Targets in Multiple Availability Zones:

Amazon EFS provides a scalable and highly available network file system that supports the NFS protocol.

EFS is ideal for Linux instances as it allows multiple instances to read and write data concurrently.

Setting up EFS with targets in multiple Availability Zones ensures high availability and durability.

Reference: Amazon EFS Overview

Create an Amazon FSx for NetApp ONTAP Multi-AZ File System:

Amazon FSx for NetApp ONTAP offers a fully managed file storage solution that supports both SMB for Windows and NFS for Linux.

The Multi-AZ deployment ensures high availability and durability, providing sub-millisecond latencies suitable for the application's performance requirements.

Reference: Amazon FSx for NetApp ONTAP

Update the User Data for Each Application's Launch Template to Mount the File System:

Updating the user data in the launch template ensures that every new instance launched by the Auto Scaling group will automatically mount the appropriate file system.

This step is necessary to ensure that all instances can access the shared storage without manual intervention.

Example user data for mounting EFS (Linux)

```
#!/bin/bash
```

```
sudo yum install -y amazon-efs-utils
```

```
sudo mount -t efs fs-12345678:/ /mnt/efs
```

Example user data for mounting FSx (Windows):

By implementing these steps, the company can provide a durable storage solution with sub-millisecond latencies that supports both SMB and NFS protocols, meeting the requirements for both Windows and Linux instances.

References:

Mounting EFS File Systems

Mounting Amazon FSx File Systems

NEW QUESTION: 171

A company is migrating its on-premises application to AWS. The application consists of a web front-end and a database. The web front-end is currently hosted on Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance. The database is currently hosted on an Amazon EC2 instance running MySQL. The company wants to improve the availability and scalability of the application. Which of the following architectures is the most appropriate?

- A. EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, Amazon EC2 instances running MySQL, and Amazon Aurora Serverless v2.
- B. EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, Amazon EC2 instances running MySQL, and Amazon DynamoDB.
- C. EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, Amazon Lambda functions, Amazon EC2 instances running MySQL, and Amazon Aurora Serverless v2.
- D. EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, Amazon Lambda functions, Amazon EC2 instances running MySQL, and Amazon DynamoDB.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 172

A company is migrating its on-premises application to AWS. The application consists of a web front-end and a database. The web front-end is currently hosted on Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance. The database is currently hosted on an Amazon EC2 instance running MySQL. The company wants to improve the availability and scalability of the application. Which of the following architectures is the most appropriate?

- A. Amazon S3, Amazon Route 53, Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, and Amazon DynamoDB.
- B. Amazon S3, Amazon Route 53, Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, and Amazon Aurora Serverless v2.
- C. Amazon S3, Amazon Route 53, Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, and Amazon DynamoDB.
- D. Amazon S3, Amazon Route 53, Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, and Amazon Aurora Serverless v2.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 173

A company is migrating its on-premises application to AWS. The application consists of a web front-end and a database. The web front-end is currently hosted on Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance. The database is currently hosted on an Amazon EC2 instance running MySQL. The company wants to improve the availability and scalability of the application. Which of the following architectures is the most appropriate?

- A. AWS Trusted Advisor, Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, Amazon CloudFormation StackSets, and Amazon DynamoDB.
- B. Amazon CloudFormation StackSets, Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, Amazon CloudFormation StackSets, and Amazon DynamoDB.
- C. Amazon CloudFormation StackSets, Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, Amazon CloudFormation StackSets, and Amazon DynamoDB.
- D. Amazon CloudFormation StackSets, Amazon EC2 instances behind an Amazon Elastic Load Balancing (ALB) instance, Amazon Service Catalog, and Amazon DynamoDB.

Answer: (SHOW ANSWER)

service catalog uses stacksets and can enforce tag and restrict resources AWS Customer case with tag enforcement
<https://aws.amazon.com/ko/blogs/apn/enforce-centralized-tag-compliance-using-aws-service-catalog-amazon-dynamodb-aws-lambda->

and-amazon-cloudwatch-events/ And Youtube video showing how to restrict resources per user with portfolio

<https://www.youtube.com/watch?v=LzvhTcqyog>

NEW QUESTION: 174

You are using Amazon Elastic Container Registry (ECR) to store Docker images for your Java application. Your application is built using CodePipeline. You want to ensure that the application is built and deployed using the latest version of the Docker image. Which of the following configurations will ensure that the application is built and deployed using the latest version of the Docker image?

Which of the following configurations will ensure that the application is built and deployed using the latest version of the Docker image?

A. Inspector is used to scan for vulnerabilities in the Docker image. EventBridge is used to trigger a Lambda function that updates the Docker image in ECR.

B. Amazon ECR is configured with Amazon Inspector for vulnerability scanning. Amazon S3 is used to store the Docker image. Athena is used to query the Docker image metadata.

C. Amazon ECR is configured with Amazon Detective for vulnerability scanning. Amazon Lambda is used to update the Docker image in ECR.

D. Amazon ECR is configured with Amazon Inspector for vulnerability scanning. Amazon EventBridge is used to trigger a Lambda function that updates the Docker image in ECR.

Answer: D (LEAVE A REPLY)

* Amazon ECR enhanced scanning uses Amazon Inspector for vulnerability detection.

* EventBridge can capture Inspector scan findings.

* Lambda can process scan findings and reject manual approval if critical vulnerabilities exist.

* Options A and C use incorrect or less integrated services (basic scanning or Detective).

* Option B adds unnecessary complexity with SBOM and Athena.

References:

Amazon ECR Image Scanning

Integrating ECR Scanning with CodePipeline

NEW QUESTION: 175

You are using Docker to build and run your application. Your application is built using CodeBuild. You want to ensure that the application is built and run using the latest version of the Docker image. Which of the following configurations will ensure that the application is built and run using the latest version of the Docker image?

DevOps is used to manage the Docker image. Amazon S3 is used to store the Docker image. Amazon Elastic Container Registry (Amazon ECR) is used to store the Docker image. CodeBuild is used to build the Docker image. AWS IAM is used to manage the Docker image.

CodeBuild is used to build the Docker image. Amazon S3 is used to store the Docker image. Amazon Elastic Container Registry (Amazon ECR) is used to store the Docker image. DevOps is used to manage the Docker image.

CodeBuild is used to build the Docker image. Amazon S3 is used to store the Docker image. Amazon Elastic Container Registry (Amazon ECR) is used to store the Docker image. DevOps is used to manage the Docker image.

CodeBuild is used to build the Docker image. Amazon S3 is used to store the Docker image. Amazon Elastic Container Registry (Amazon ECR) is used to store the Docker image. DevOps is used to manage the Docker image.

A. CodeBuild is used to build the Docker image. Amazon S3 is used to store the Docker image. Amazon Elastic Container Registry (Amazon ECR) is used to store the Docker image. DevOps is used to manage the Docker image.

B. CodeBuild is used to build the Docker image. Amazon S3 is used to store the Docker image. Amazon Elastic Container Registry (Amazon ECR) is used to store the Docker image. DevOps is used to manage the Docker image.

C. CodeBuild is used to build the Docker image. Amazon S3 is used to store the Docker image. Amazon Elastic Container Registry (Amazon ECR) is used to store the Docker image. DevOps is used to manage the Docker image.

D. CodeBuild is used to build the Docker image. Amazon S3 is used to store the Docker image. Amazon Elastic Container Registry (Amazon ECR) is used to store the Docker image. DevOps is used to manage the Docker image.

Answer: A (LEAVE A REPLY)

Update the buildspec.yml file to log in to the ECR repository by using the `aws ecr get-login-password` AWS CLI command to obtain an authentication token. Update the docker login command to use the authentication token to access the ECR repository. This is the correct solution. The `aws ecr get-login-password` AWS CLI command retrieves and displays an authentication token that can be used to log in to an ECR repository. The docker login command can use this token as a password to authenticate with the ECR repository. This way, the CodeBuild project can push and pull images from the ECR repository without any errors. For more information, see [Using Amazon ECR with the AWS CLI and get-login-password](#).

NEW QUESTION: 176

A Python CodeArtifact repository is configured with the following permissions policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "codeartifact:DescribePackageVersion", "codeartifact:DescribeRepository",
        "codeartifact:GetPackageVersionReadme", "codeartifact:GetRepositoryEndpoint", "codeartifact:
        ListPackageVersionAssets", "codeartifact: ListPackageVersionDependency", "codeartifact:
        ListPackageVersions", "codeartifact :ListPackages",
        "codeartifact: ReadFromRepository"
      ],
      "Resource": "*"
    }
  ],
  "Principal": {
    "AWS": "*"
  }
}

```

The repository is used to store Python packages. The repository is located in a VPC. The repository is used to store Python packages. The repository is used to store Python packages.

The repository is used to store Python packages. The repository is used to store Python packages. The repository is used to store Python packages.

DevOps wants to use the repository to store Python packages. Which of the following is the correct solution? (Select TWO.)

- A. Amazon S3 bucket and CodeBuild project with the appropriate IAM role.
- B. IAM role with the appropriate permissions and CodeBuild project with the appropriate IAM role.
- C. AWS Resource Access Manager(AWS RAM) role and CodeArtifact repository.

* cdk diff checks for changes but is not a unit test and may not catch logical errors.

References:

Testing AWS CDK Applications

CodeBuild Buildspec OnFailure

NEW QUESTION: 180

Which of the following are supported by AWS Organizations? (Select two.)
A. AWS Control Tower
B. AWS Center for Internet Security (CIS) Benchmarks to AWS Foundations

C. AWS Security Hub
D. AWS Security Hub Security Hub
E. AWS Security Hub Security Hub

Which of the following are supported by AWS Organizations? (Select three.)

A. AWS Security Hub
B. AWS Security Hub
C. AWS Security Hub
D. AWS Security Hub
E. AWS Security Hub

B. AWS Security Hub
C. AWS Security Hub
D. AWS Security Hub

C. AWS IAM (AWS Single Sign-On) CreateAccountAssignment API
D. AWS Security Hub SCP

D. AWS Security Hub SCP

E. AWS Security Hub

F. CreateManagedAccount Amazon EventBridge Security Hub CreateMembers API Security Hub AWS Lambda Lambda EventBridge

Answer: A,C,E (LEAVE A REPLY)

<https://docs.aws.amazon.com/securityhub/latest/userguide/accounts-orgs-auto-enable.html>

NEW QUESTION: 181

Which of the following are supported by AWS Organizations? (Select two.)
A. Amazon API Gateway API
B. Amazon S3 SDK
C. Amazon CloudFront SDK
D. DevOps API SDK

Which of the following are supported by AWS Organizations? (Select two.)

A. API Gateway SDK
B. API Gateway SDK
C. API Gateway SDK
D. API Gateway SDK

B. API Gateway SDK
C. API Gateway SDK
D. API Gateway SDK

C. aws apigateway UpdateStage triggers an Amazon EventBridge event. AWS Lambda function invokes API Gateway SDK. SDK S3 bucket CloudFront API invalidation. SDK.

D. Create Amazon EventBridge event. aws apigateway API SDK AWS Lambda function. Gateway SDK S3 bucket S3 API SDK bucket.

Answer: A (LEAVE A REPLY)

This solution would allow the company to automate the process of updating the SDK and making it available to web clients. By adding a CodePipeline action immediately after the deployment stage of the API, the Lambda function will be invoked automatically each time the API is updated. The Lambda function should be able to download the new SDK from API Gateway, upload it to the S3 bucket and also create a CloudFront invalidation for the SDK path so that the latest version of the SDK is available for the web clients. This is the most straight forward solution and it will meet the requirements.

DOP-C02-KR DumpTop DOP-C02-KR! DumpTop **DOP-C02-KR** DumpTop DOP-C02-KR, DumpTop DOP-C02-KR. <https://www.dumptop.com/Amazon/DOP-C02-KR-dump.html> (439 Q&As Dumps, 30%OFF Special Discount: **KrDump**)

NEW QUESTION: 182

DevOps team uses Amazon ECR and AWS CodePipeline. They want to ensure that new executions wait for the previous execution to finish, preventing overlapping runs. They also want to listen for Git tag pushes. An EventBridge rule can detect new image pushes to ECR and start the deployment pipeline, integrating the build and deploy pipelines effectively.

- A. QUEUED mode ensures that new executions wait for the previous execution to finish, preventing overlapping runs. Adding a trigger filter to listen for Git tag triggers the pipeline only on new tag pushes.
- B. SUPERSEDED mode cancels the previous run when a new one starts, which is not desired here.
- C. SUPERSEDED mode ensures that new executions wait for the previous execution to finish, preventing overlapping runs. Adding a trigger filter to listen for Git tag triggers the pipeline only on new tag pushes.
- D. QUEUED mode ensures that new executions wait for the previous execution to finish, preventing overlapping runs. Adding a trigger filter to listen for Git tag triggers the pipeline only on new tag pushes.

Answer: (SHOW ANSWER)

- * CodePipeline V2 with QUEUED mode ensures that new executions wait for the previous execution to finish, preventing overlapping runs.
- * Adding a trigger filter to listen for Git tag triggers the pipeline only on new tag pushes.
- * An EventBridge rule can detect new image pushes to ECR and start the deployment pipeline, integrating the build and deploy pipelines effectively.
- * SUPERSEDED mode cancels the previous run when a new one starts, which is not desired here.
- * Using branches instead of tags would trigger on all commits, not just releases.

References:

AWS CodePipeline V2 Triggers

AWS CodePipeline Execution Modes

NEW QUESTION: 183

□□□ ALB(Application Load Balancer) □□ □□ Amazon EC2 □□□□□ □ □□□□□□□ □□□□□. □□□ AWS CodeCommit □ □□□□□ □□□□□□ □□□ □□□□□. □□□ □□ □□□□ □□□□ AWS Lambda □□□ AWS CodeBuild □□□□□ □□□□ □. CodeBuild □□□□□ □□□ □□□□□, □□□□□ □□□ AWS CodeArtifact□ □□□□, AWS □□□ □□□ Run Command□ □□ □□ □□□□ □□□ EC2 □□□□□ □□□□□.

□□ □□□□□ □□, □□□□□ □□□ □□ □□□ □□□□ □□ EC2 □□□□□, □□□□□ □ □□□□□ □□□□□□□.

□□ □□□□ □□ □□□□ □□□□□ DevOps □□□□□ □□□□□ □□ □□ □□□ □□□□□? (2□□ □□□□□.)

A. AWS CodePipeline□□ CodeCommit □□□□□□ □□ □□□□ □□□□ □□□□□□ □□□□□□. □□□□□□□□ □□□□ □□ □□□ □□ CodeBuild □□□□□ □□□ □□□□ □□□□□ □□□ □□□□□. □□□□□□□□ CodeBuild □□□□□ □□ □□□□□□ AWS CodeDeploy □□□ □□□□□□.

B. AWS CodePipeline□□ CodeCommit □□□□□□ □□ □□□□ □□□□ □□□□□□ □□□□□□. CodeBuild □□□□□ □□□□ □□□□□□□□ □□□□ □□□□□ □□□ □□□□□ □□□ □□□□□□. □□□□□□□□ CodeBuild □□□□□ □□ □□□□□□ AWS CodeDeploy □□□ □□□□□□.

C. AWS CodeDeploy □□□□□□□□ □□ □□□ □□□□ □□□□ □□□ EC2 □□□□□□ □□□□□□. □□ □□□ □□ ALB□ □□ □□□□.

D. Systems Manager □□ AWS CodeDeploy□ □□□□□ □□, □□□□ □ □□ □□□ □□□□ □□ Lambda □□□ □□□□□□.

E. Amazon S3 □□□ □□□□□□. CodeArtifact □□ S3 □□□ □□□□ □□□□□□ CodeBuild □□□□□□ □□□□□□. CodeDeploy□ □□ □□□ □□□□□ □□□□□□ EC2 □□□□□□ □□□□□□.

Answer: A,C (LEAVE A REPLY)

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 configuration. 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 traffic. 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

- 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

DOP-C02-KR ☐☐ ☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop ☐☐ ☐☐☐☐ ☐☐☐ DOP-C02-KR ☐☐! DumpTop ☐ ☐☐ **DOP-C02-KR** ☐☐ ☐☐☐ ☐☐☐☐☐☐, DumpTop DOP-C02-KR ☐☐ ☐☐☐ ☐☐☐☐☐☐☐☐☐ ☐☐☐ ☐☐☐☐☐☐☐☐. ☐☐☐☐☐ ☐☐☐☐☐☐☐☐☐ ☐☐☐☐☐☐☐☐☐. <https://www.dumptop.com/Amazon/DOP-C02-KR-dump.html> (439 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)