Exam

SAP-C02

AWS Certified Solutions Architect – Professional

Free demo	80 sample questions
Full exam	681 questions with answers and explanations
Contents	Questions with answers and explanations

Exam domains

Domain	Exam weight	Questions
D1 Design Solutions for Organizational Complexity	12%	20
D2 Design for New Solutions	31%	23
D3 Continuous Improvement for Existing Solutions	29%	25
D4 Accelerate Workload Migration and Modernization	28%	12

This is a free demo from Dumpsgate: 80 of the exam's 681 questions. **Get the full SAP-C02 exam** with every question, answer and explanation.

Question 1

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A company runs applications in multiple AWS accounts in a single AWS Region. The accounts are members of an organization in AWS Organizations. The company is implementing security controls to manage cross-account access patterns in the company's AWS environment. The company wants to identify resources that might be unintentionally shared with external entities. The company wants to implement continuous monitoring for resource-based policies to detect external access. The company needs a solution that enforces multi-factor authentication (MFA) to assume cross-account IAM roles.

Which solution will meet these requirements?

A. Configure AWS IAM Access Analyzer to use the organization as the zone of trust. Configure an Amazon EventBridge rule to send a notification when IAM Access Analyzer generates new findings. Add the aws:MultiFactorAuthPresent condition key to the trust policies of employee IAM roles.

B. Configure AWS Config rules to monitor resource-based policies. Use IAM permissions boundaries to limit cross-account access. Add the aws:TokenIssueTime condition key to cross-account IAM roles.

C. Enable AWS IAM Access Analyzer. Implement ExternalId in cross-account IAM roles. Create a resource-based policy that uses a wildcard principal to limit access based on the source IP.

D. Create an AWS Security Hub integration with AWS IAM Access Analyzer. Implement federated identity management for all cross-account IAM role access. Enforce MFA by using a third-party identity provider (IdP).

Answer: A. Configure AWS IAM Access Analyzer to use the organization as the zone of trust. Configure an Amazon EventBridge rule to send a notification when IAM Access Analyzer generates new findings. Add the aws:MultiFactorAuthPresent condition key to the trust policies of employee IAM roles.

Explanation

A is the only option that addresses all three requirements: an organization-level IAM Access Analyzer can identify resources whose policies grant access outside the organization, and EventBridge can notify the company about new findings. Adding the `aws:MultiFactorAuthPresent` condition to the relevant role trust policies requires MFA when employees assume those roles. Together, these controls provide external-sharing detection, ongoing notification, and MFA enforcement for cross-account role assumption.

https://docs.aws.amazon.com/IAM/latest/UserGuide/access-analyzer-findings.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access-analyzer-eventbridge.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access-analyzer-concepts.html
https://docs.aws.amazon.com/eventbridge/latest/ref/events-ref-access-analyzer.html

✓ **A** Configure AWS IAM Access Analyzer to use the organization as the zone of trust. Configure an Amazon EventBridge rule to send a notification when IAM Access Analyzer generates new findings. Add the aws:MultiFactorAuthPresent condition key to the trust policies of employee IAM roles.
An organization-level analyzer uses the organization as its zone of trust to find resources shared with external entities. EventBridge can route new Access Analyzer findings for notification, supporting continuous monitoring. The `aws:MultiFactorAuthPresent` condition in the role trust policies requires MFA for the specified role assumptions.

https://docs.aws.amazon.com/IAM/latest/UserGuide/access-analyzer-findings.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access-analyzer-eventbridge.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access-analyzer-concepts.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/eventbridge/latest/ref/events-ref-access-analyzer.html

✗ **B** Configure AWS Config rules to monitor resource-based policies. Use IAM permissions boundaries to limit cross-account access. Add the aws:TokenIssueTime condition key to cross-account IAM roles.
This does not use IAM Access Analyzer to identify external access in resource-based policies, and permissions boundaries do not by themselves provide that external-sharing detection. The `aws:TokenIssueTime` condition key does not enforce MFA when assuming a role.

✗ **C** Enable AWS IAM Access Analyzer. Implement ExternalId in cross-account IAM roles. Create a resource-based policy that uses a wildcard principal to limit access based on the source IP.
Although Access Analyzer is included, the option does not configure the organization as the zone of trust or set up notifications for new findings. ExternalId is not an MFA requirement, and a wildcard principal constrained by source IP does not satisfy the requirement to enforce MFA for cross-account role assumption.

https://docs.aws.amazon.com/IAM/latest/UserGuide/access-analyzer-concepts.html

✗ **D** Create an AWS Security Hub integration with AWS IAM Access Analyzer. Implement federated identity management for all cross-account IAM role access. Enforce MFA by using a third-party identity provider (IdP).
An IdP can enforce MFA for federated access, but this option does not specify configuring IAM Access Analyzer with the organization as its zone of trust or notifying on new findings. A Security Hub integration alone does not establish the specified Access Analyzer monitoring and notification setup.

https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-internal-providers.html

Dumpsgate

Question 2

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A financial company uses AWS Control Tower to govern multiple AWS accounts. The company must comply with data residency regulations. The regulations require customer data and backups to remain only in specific approved AWS Regions. Additionally, the company must retain control of the cryptographic root of trust. The company must use encryption keys in hardware security modules (HSMs) that the company operates in-country. The company needs a preventive control that blocks resource creation outside the approved Regions.

Which solution will meet these requirements with the LEAST operational overhead?

A. Use an AWS KMS external key store (XKS) backed by the company's in-country HSMs for encryption. Apply an AWS Control Tower Region deny control to the relevant OUs to deny actions outside the approved Regions.

B. Use AWS KMS multi-Region keys for all encryption. Use AWS Config rules to detect resources that are created outside the approved Regions.

C. Use an AWS KMS custom key store backed by an AWS CloudHSM cluster deployed in one approved Region. Apply separate IAM policies in each account. Configure the policies to deny actions outside the approved Regions by using the aws:RequestedRegion condition key.

D. Use standard AWS KMS customer managed keys for encryption. Apply an AWS Control Tower Region deny control to the relevant OUs to deny actions outside the approved Regions.

Answer: A. Use an AWS KMS external key store (XKS) backed by the company's in-country HSMs for encryption. Apply an AWS Control Tower Region deny control to the relevant OUs to deny actions outside the approved Regions.

Dumpsgate

Question 3

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A company has more than 20 application development teams. All the teams use AWS CloudFormation to deploy application resources in the company's production AWS account. The production account is a member of an organization in AWS Organizations that has all features enabled.

The company must ensure that all resources that the development teams provision match a set of predefined rules. The company needs a solution to prevent users from provisioning noncompliant resources through CloudFormation.

Which solution will meet these requirements with the LEAST operational overhead?

A. Implement each configuration rule as a separate AWS Config custom AWS Lambda rule. Deploy the Lambda functions to the company's production account. Implement SCPs to prevent changes to the Lambda functions and to deny any operation that could disable AWS Config or its rules or triggers in the production account.

B. Use CloudFormation Hooks to implement the configuration rules. Use a separate AWS Lambda function to enforce each rule. Use CloudFormation StackSets to deploy and register Hooks across all member accounts in the organization. Use SCPs and IAM policies to prevent Hook modifications and Lambda function changes in member accounts.

C. Implement the configuration rules as an AWS Lambda function. Configure a CloudFormation macro definition to make the Lambda function available to CloudFormation. Configure the function to check the CloudFormation template against the configuration rules. Use the macro in all CloudFormation deployment templates.

D. Implement the configuration rules by using AWS CloudFormation Guard rules. Use the cfn-guard-lambda command to install Guard as an AWS Lambda function. Build and run Guard by invoking the Lambda function. Provide the CloudFormation template and the rule set as inputs to the function.

Answer: B. Use CloudFormation Hooks to implement the configuration rules. Use a separate AWS Lambda function to enforce each rule. Use CloudFormation StackSets to deploy and register Hooks across all member accounts in the organization. Use SCPs and IAM policies to prevent Hook modifications and Lambda function changes in member accounts.

Explanation

CloudFormation Hooks can evaluate resources during CloudFormation operations and block operations that fail the configured checks. Registering the Hooks across the organization with StackSets gives the company a way to apply the same rules across member accounts, rather than relying on each development team to include a macro or run a separate validation step. Protecting the Hooks and their Lambda functions from modification helps keep the controls in place. This meets the prevention requirement with less ongoing effort than managing independent account-level rules or relying on teams to invoke validation themselves.

https://docs.aws.amazon.com/cloudformation-cli/latest/hooks-userguide/registering-hooks.html
https://docs.aws.amazon.com/cloudformation-cli/latest/hooks-userguide/lambda-hooks-create-lambda-function.html
https://docs.aws.amazon.com/cloudformation-cli/latest/hooks-userguide/hooks-publishing.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html

✗ **A** Implement each configuration rule as a separate AWS Config custom AWS Lambda rule. Deploy the

Free demo · full exam: 681 questions

Lambda functions to the company's production account. Implement SCPs to prevent changes to the Lambda functions and to deny any operation that could disable AWS Config or its rules or triggers in the production account.
AWS Config rules are primarily for evaluating resource configuration, so this approach detects noncompliance rather than reliably blocking a CloudFormation operation before resources are provisioned. It also requires separate rules and functions to be managed in the production account, adding operational work. Protecting the configuration with SCPs does not turn the evaluations into CloudFormation deployment-time enforcement.

https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_components.html

✓ **B** Use CloudFormation Hooks to implement the configuration rules. Use a separate AWS Lambda function to enforce each rule. Use CloudFormation StackSets to deploy and register Hooks across all member accounts in the organization. Use SCPs and IAM policies to prevent Hook modifications and Lambda function changes in member accounts.
CloudFormation Hooks can evaluate CloudFormation operations and prevent operations from proceeding when resources fail the configured checks. StackSets can distribute the Hooks across the organization's member accounts, while SCPs and IAM policies help prevent teams from changing or bypassing the controls. This applies the rules to CloudFormation deployments centrally instead of depending on each team to perform a separate validation step.

https://docs.aws.amazon.com/cloudformation-cli/latest/hooks-userguide/registering-hooks.html
https://docs.aws.amazon.com/cloudformation-cli/latest/hooks-userguide/lambda-hooks-create-lambda-function.html
https://docs.aws.amazon.com/cloudformation-cli/latest/hooks-userguide/hooks-publishing.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html

✗ **C** Implement the configuration rules as an AWS Lambda function. Configure a CloudFormation macro definition to make the Lambda function available to CloudFormation. Configure the function to check the CloudFormation template against the configuration rules. Use the macro in all CloudFormation deployment templates.
A macro is invoked when a template includes the macro's transform; it does not automatically enforce checks on every CloudFormation deployment. The approach therefore depends on all teams including the macro in every applicable template, and a team that omits it can bypass the validation. That is less reliable than centrally registered Hooks for the required prevention.

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-macros.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-macros-overview.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-macros-author.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/macros-example.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/TemplateReference/transform-reference.html

✗ **D** Implement the configuration rules by using AWS CloudFormation Guard rules. Use the cfn-guard-lambda command to install Guard as an AWS Lambda function. Build and run Guard by invoking the Lambda function. Provide the CloudFormation template and the rule set as inputs to the function.
This runs Guard only when the Lambda function is invoked with a template and rule set; it does not by itself make CloudFormation invoke the validation or block noncompliant deployments. The company would need an additional enforced integration to ensure every team's CloudFormation operation goes through that function. As described, this does not meet the requirement to prevent users from provisioning noncompliant resources.

Free demo · full exam: 681 questions

Question 4

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A multinational media company is building a video processing platform on AWS. The platform is deployed across multiple AWS accounts in an organization in AWS Organizations. Customers upload videos through Amazon API Gateway. AWS Lambda functions process the videos. The company stores processed media and metadata in Amazon S3 and Amazon DynamoDB. The company serves customers globally.

Regulatory requirements state that the company must process and store European Union (EU) customer content only within AWS Regions in the EU. The company must prevent workloads outside the EU sovereignty boundary from processing EU customer content. The company needs preventive governance controls so that teams cannot accidentally deploy EU workloads or infrastructure outside EU Regions.

Which solution will meet these requirements?

A. Deploy separate EU and non-EU processing stacks by using API Gateway and Lambda. Use Amazon Route 53 geolocation routing to send EU uploads to EU endpoints. Store EU media in Amazon S3 in EU Regions. Store EU metadata in DynamoDB in EU Regions. Enforce Region restrictions by using Organizations SCPs.

B. Deploy a single global API Gateway endpoint that invokes Lambda functions in multiple Regions. Use Amazon Route 53 geolocation routing for EU uploads. Store EU media in an S3 bucket in the eu-west-1 Region. Store EU metadata in a DynamoDB table in the same Region. Restrict access by using IAM policies.

C. Deploy API Gateway and Lambda workloads only in EU Regions for all customer uploads. Use Amazon Route 53 geolocation routing for EU uploads. Store EU media in Amazon S3 in EU Regions. Store EU metadata in DynamoDB in EU Regions. Use AWS Config rules to detect resources that are deployed outside EU Regions.

D. Deploy separate EU and non-EU processing stacks by using API Gateway and Lambda. Use Amazon Route 53 latency-based routing to send requests to the nearest endpoint. Store EU media in Amazon S3 in EU Regions. Store EU metadata in DynamoDB in EU Regions. Use AWS Control Tower controls to monitor deployments outside EU Regions.

Answer: A. Deploy separate EU and non-EU processing stacks by using API Gateway and Lambda. Use Amazon Route 53 geolocation routing to send EU uploads to EU endpoints. Store EU media in Amazon S3 in EU Regions. Store EU metadata in DynamoDB in EU Regions. Enforce Region restrictions by using Organizations SCPs.

Explanation

A separates EU and non-EU processing stacks and routes EU uploads to EU endpoints, so EU content can be processed and stored in EU Regions. Storing the media and metadata in EU Regions addresses the data-location requirement. Organizations SCPs provide preventive Region restrictions across accounts, helping prevent teams from deploying EU workloads or infrastructure outside the allowed Regions.

https://docs.aws.amazon.com/controltower/latest/userguide/region-deny.html
https://docs.aws.amazon.com/controltower/latest/controlreference/ou-region-deny.html

✓ **A** Deploy separate EU and non-EU processing stacks by using API Gateway and Lambda. Use Amazon Route 53 geolocation routing to send EU uploads to EU endpoints. Store EU media in Amazon S3 in EU Regions. Store EU metadata in DynamoDB in EU Regions. Enforce Region restrictions by using Organizations SCPs.

Free demo · full exam: 681 questions

Separate processing stacks and geolocation routing direct EU uploads to EU endpoints, while the S3 and DynamoDB placement keeps EU content in EU Regions. SCPs can impose Region restrictions across the organization's accounts as a preventive governance control, rather than relying only on detecting an out-of-Region deployment afterward.

https://docs.aws.amazon.com/controltower/latest/userguide/region-deny.html
https://docs.aws.amazon.com/controltower/latest/controlreference/ou-region-deny.html

✗ **B** Deploy a single global API Gateway endpoint that invokes Lambda functions in multiple Regions. Use Amazon Route 53 geolocation routing for EU uploads. Store EU media in an S3 bucket in the eu-west-1 Region. Store EU metadata in a DynamoDB table in the same Region. Restrict access by using IAM policies.
Although the data stores are in an EU Region, the global endpoint can invoke Lambda functions in multiple Regions, so EU content could be processed outside the sovereignty boundary. IAM policies alone do not provide the stated organization-wide preventive Region governance that SCPs can enforce.

https://docs.aws.amazon.com/controltower/latest/userguide/region-deny.html
https://docs.aws.amazon.com/controltower/latest/controlreference/ou-region-deny.html

✗ **C** Deploy API Gateway and Lambda workloads only in EU Regions for all customer uploads. Use Amazon Route 53 geolocation routing for EU uploads. Store EU media in Amazon S3 in EU Regions. Store EU metadata in DynamoDB in EU Regions. Use AWS Config rules to detect resources that are deployed outside EU Regions.
AWS Config rules can evaluate resources and identify out-of-Region deployments, but detection does not itself prevent teams from creating them. The requirement calls for preventive governance controls, such as Region restrictions enforced through SCPs.

https://docs.aws.amazon.com/us_en/config/latest/developerguide/evaluating-your-resources.html
https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config.html

✗ **D** Deploy separate EU and non-EU processing stacks by using API Gateway and Lambda. Use Amazon Route 53 latency-based routing to send requests to the nearest endpoint. Store EU media in Amazon S3 in EU Regions. Store EU metadata in DynamoDB in EU Regions. Use AWS Control Tower controls to monitor deployments outside EU Regions.
Latency-based routing chooses a nearby endpoint, not necessarily an endpoint inside the EU, so it could send EU uploads to processing outside the sovereignty boundary. The option also describes monitoring deployments, rather than applying a preventive Region restriction to block them.

https://docs.aws.amazon.com/controltower/latest/userguide/how-controls-work.html
https://docs.aws.amazon.com/controltower/latest/controlreference/control-behavior.html

Free demo · full exam: 681 questions

Question 5

Multiple choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A company needs to apply consistent custom security audit rules across its organization in AWS Organizations. Accounts in the organization use a single AWS Region.

The company uses AWS Config in each account with four custom rules and five managed rules. The rules must be deployed to all accounts and must not be editable by member accounts. The company wants to easily update rules in the future. All required permissions are already in place.

Which combination of steps will meet these requirements with the LEAST operational effort? (Select TWO.)

A. Call the PutOrganizationConformancePack API operation from a delegated administrator account to deploy the conformance pack to the entire organization.

B. Create an AWS Config conformance pack that includes four custom policy rules.

C. Create an AWS Config conformance pack that includes four custom policy rules and five AWS managed rules.

D. Deploy five AWS organizational rules and the conformance pack from a delegated administrator account to the entire organization.

E. Use AWS CloudFormation StackSets to deploy the conformance pack with termination protection.

Answer

- A. **Call the PutOrganizationConformancePack API operation from a delegated administrator account to deploy the conformance pack to the entire organization.**
- C. **Create an AWS Config conformance pack that includes four custom policy rules and five AWS managed rules.**

Explanation

An AWS Config conformance pack can bundle the four custom policy rules and five managed rules, making them straightforward to update together. Deploying that pack with the organization-level API from a delegated administrator applies it across the organization and centralizes management, rather than requiring separate deployments in every account. This meets the requirements with less operational effort than combining separate organizational rules or using StackSets.

https://docs.aws.amazon.com/config/latest/developerguide/custom-conformance-pack.html
https://docs.aws.amazon.com/config/latest/developerguide/conformancepack-sample-templates.html
https://docs.aws.amazon.com/cli/latest/reference/configservice/put-organization-conformance-pack.html

✓ **A** Call the PutOrganizationConformancePack API operation from a delegated administrator account to deploy the conformance pack to the entire organization.
This deploys the conformance pack organization-wide from a central account, avoiding separate deployments in member accounts. The organization-level operation supports managing the deployment centrally, which fits the requirement that member accounts must not edit the rules.

https://docs.aws.amazon.com/cli/latest/reference/configservice/put-organization-conformance-pack.html
https://docs.aws.amazon.com/config/latest/developerguide/cpack-prerequisites.html

✗ **B** Create an AWS Config conformance pack that includes four custom policy rules.
This pack contains the four custom rules but omits the five managed rules required by the scenario. The conformance pack must include both types so that all nine rules can be deployed and updated together.

https://docs.aws.amazon.com/config/latest/developerguide/custom-conformance-pack.html
https://docs.aws.amazon.com/config/latest/developerguide/conformancepack-sample-templates.html

✓ **C** Create an AWS Config conformance pack that includes four custom policy rules and five AWS

Free demo · full exam: 681 questions

managed rules.
A conformance pack can group custom policy rules and AWS managed rules into one deployable set. Including all nine rules makes future updates easier because the company can maintain them together.

https://docs.aws.amazon.com/config/latest/developerguide/custom-conformance-pack.html
https://docs.aws.amazon.com/config/latest/developerguide/conformancepack-sample-templates.html

✗ **D** Deploy five AWS organizational rules and the conformance pack from a delegated administrator account to the entire organization.
The five managed rules can be included in the same conformance pack as the four custom rules, so deploying them separately as organizational rules adds operational work. This approach is not the least-effort combination for managing and updating all required rules together.

https://docs.aws.amazon.com/config/latest/developerguide/custom-conformance-pack.html
https://docs.aws.amazon.com/config/latest/developerguide/conformancepack-sample-templates.html
https://docs.aws.amazon.com/cli/latest/reference/configservice/put-organization-conformance-pack.html

✗ **E** Use AWS CloudFormation StackSets to deploy the conformance pack with termination protection.
StackSets are not needed when AWS Config provides an organization-level conformance-pack deployment mechanism. This adds a separate deployment method and termination protection does not itself ensure that member accounts cannot edit the rules.

https://docs.aws.amazon.com/cli/latest/reference/configservice/put-organization-conformance-pack.html

Dumpsgate

Question 6

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A global healthcare analytics company runs a regulated workload on AWS across dozens of AWS accounts. The company uses an organization in AWS Organizations to manage the accounts. The company must regularly provide external auditors with evidence that specific security controls are implemented and continuously enforced. The security controls include encryption requirements for storage services, centralized logging configurations, and restrictions on public network access.

The company wants an automated solution that continuously collects evidence that shows that the controls are implemented across accounts. The solution must preserve historical evidence for specified time periods. The solution must also generate reports for the auditors that are mapped to specific regulatory frameworks. The company does not want to build custom evidence collection pipelines.

Which solution will meet these requirements with the LEAST operational overhead?

A. Enable AWS CloudTrail organization trails across all accounts and deliver the logs to a centralized Amazon S3 bucket in an audit account. Use Amazon Athena to query the centralized logs and build scheduled reports that auditors can review for evidence of the required security controls.

B. Enable AWS Security Hub across all accounts. Designate a delegated administrator account to aggregate security findings across the organization. Export Security Hub findings and compliance check results to Amazon S3. Generate periodic compliance reports for auditors.

C. Enable AWS Config rules in each account to evaluate the required security controls. Deliver configuration snapshots to a centralized Amazon S3 bucket. Use AWS Lambda functions in an audit account to periodically analyze the snapshots and generate compliance evidence reports for auditors.

D. Enable AWS Config rules across all accounts to evaluate the required security controls. Use an AWS Config aggregator in an audit account to centralize configuration data. Deploy AWS Config conformance packs that are aligned to the required regulatory frameworks. Use AWS Audit Manager to collect evidence and generate audit reports.

Answer: D. Enable AWS Config rules across all accounts to evaluate the required security controls. Use an AWS Config aggregator in an audit account to centralize configuration data. Deploy AWS Config conformance packs that are aligned to the required regulatory frameworks. Use AWS Audit Manager to collect evidence and generate audit reports.

Explanation

Option D combines organization-wide AWS Config evaluation with conformance packs aligned to regulatory frameworks, so the company can assess controls across its AWS accounts without building its own evidence pipeline. AWS Audit Manager collects evidence for assessments and generates reports that auditors can review, while supporting evidence retention for the assessment period. This best meets the requirements for continuous evidence collection, framework-mapped reports, and low operational overhead.

https://docs.aws.amazon.com/config/latest/developerguide/conformance-pack-organization-apis.html
https://docs.aws.amazon.com/audit-manager/latest/userguide/generate-assessment-report-generation-steps.html
https://docs.aws.amazon.com/audit-manager/latest/userguide/assessment-reports.html
https://docs.aws.amazon.com/config/latest/APIReference/API_PutOrganizationConformancePack.html
https://docs.aws.amazon.com/audit-manager/latest/userguide/what-is.html

✗ **A** Enable AWS CloudTrail organization trails across all accounts and deliver the logs to a centralized

Free demo · full exam: 681 questions

Amazon S3 bucket in an audit account. Use Amazon Athena to query the centralized logs and build scheduled reports that auditors can review for evidence of the required security controls.
CloudTrail records account activity, but its logs do not by themselves establish that current configurations—such as storage encryption or public-access restrictions—comply with the controls. Querying logs and building reports would also require a custom evidence and reporting process, rather than using a managed audit-evidence solution.

✗ **B** Enable AWS Security Hub across all accounts. Designate a delegated administrator account to aggregate security findings across the organization. Export Security Hub findings and compliance check results to Amazon S3. Generate periodic compliance reports for auditors.
Security Hub can centrally aggregate findings and export them, but the proposed S3 export and report generation still leaves the company to build its own auditor-reporting workflow. It does not provide the integrated framework assessment and evidence-reporting approach offered by Audit Manager in option D.

https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-v2-findings-export.html
https://docs.aws.amazon.com/securityhub/latest/userguide/central-configuration-in-context.html
https://docs.aws.amazon.com/securityhub/latest/userguide/aggregation-central-configuration.html
https://docs.aws.amazon.com/securityhub/latest/userguide/central-configuration-intro.html

✗ **C** Enable AWS Config rules in each account to evaluate the required security controls. Deliver configuration snapshots to a centralized Amazon S3 bucket. Use AWS Lambda functions in an audit account to periodically analyze the snapshots and generate compliance evidence reports for auditors.
AWS Config can evaluate configuration state, but this option requires custom Lambda analysis and report generation. That adds the evidence-pipeline maintenance the company wants to avoid and does not provide Audit Manager's managed framework-based assessment reports.

✓ **D** Enable AWS Config rules across all accounts to evaluate the required security controls. Use an AWS Config aggregator in an audit account to centralize configuration data. Deploy AWS Config conformance packs that are aligned to the required regulatory frameworks. Use AWS Audit Manager to collect evidence and generate audit reports.
AWS Config conformance packs can be deployed across an organization to evaluate related controls, and the aggregator provides a centralized view of configuration data. Audit Manager collects assessment evidence and generates reports for auditors, avoiding custom collection and reporting code while supporting framework-aligned assessments and evidence retention.

https://docs.aws.amazon.com/config/latest/developerguide/conformance-pack-organization-apis.html
https://docs.aws.amazon.com/audit-manager/latest/userguide/generate-assessment-report-generation-steps.html
https://docs.aws.amazon.com/audit-manager/latest/userguide/assessment-reports.html
https://docs.aws.amazon.com/audit-manager/latest/userguide/examples-of-controls.html
https://docs.aws.amazon.com/config/latest/APIReference/API_PutOrganizationConformancePack.html
https://docs.aws.amazon.com/audit-manager/latest/userguide/what-is.html

Free demo · full exam: 681 questions

Question 7

Multiple choice · Domain: D1 Design Solutions for Organizational Complexity › Determine appropriate data and network protection controls.

A healthcare company is designing a patient portal. The portal must use a database to store patient dat a. The portal must encrypt all data at rest and in transit. The portal must audit all data access.

Database users must not be able to tamper with the audit trail. The company must retain the audit data for 10 years. The audit data must be immutable.

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

A. Store data in an Amazon Aurora PostgreSQL encrypted database. Enforce the use of TLS for all connections.

B. Store data in an Amazon RDS for PostgreSQL encrypted database. Enforce the use of TLS for all connections.

C. Create a database activity stream in the database cluster. Use Amazon Data Firehose to send the data to an Amazon S3 bucket. Enable S3 Object Lock in compliance mode with a 10-year retention period.

D. Install pgAudit in the database cluster. Store the data in Amazon CloudWatch Logs. Use Amazon Data Firehose to send the data to an Amazon S3 bucket. Enable S3 Object Lock in compliance mode with a 10-year retention period.

E. Install pgAudit in the database cluster. Store the data in Amazon CloudWatch Logs. Use Amazon Data Firehose to send the data to an Amazon S3 bucket. Enable S3 Object Lock in governance mode with a 10-year retention period.

Answer

- A. **Store data in an Amazon Aurora PostgreSQL encrypted database. Enforce the use of TLS for all connections.**
- C. **Create a database activity stream in the database cluster. Use Amazon Data Firehose to send the data to an Amazon S3 bucket. Enable S3 Object Lock in compliance mode with a 10-year retention period.**

Explanation

Amazon Aurora PostgreSQL can encrypt database data at rest, and TLS can protect connections in transit. A database activity stream sends activity records outside the database, while Amazon S3 Object Lock in compliance mode prevents those records from being changed or deleted during the 10-year retention period. Together, A and C address encryption, auditing, resistance to database-user tampering, and long-term immutability.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html
https://docs.aws.amazon.com/us_en/AmazonS3/latest/userguide/object-lock-managing.html
https://docs.aws.amazon.com/en_en/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.AuditLog.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.Encryption.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.SSL.html

✓ **A** Store data in an Amazon Aurora PostgreSQL encrypted database. Enforce the use of TLS for all connections.
Aurora encryption protects database data at rest, and enforcing TLS protects data in transit. This satisfies the portal's encryption requirements and pairs with the Aurora activity stream in option C.

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Security.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.Encryption.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.SSL.html

✗ **B** Store data in an Amazon RDS for PostgreSQL encrypted database. Enforce the use of TLS for all connections.
Encryption and TLS address the at-rest and in-transit requirements, but this is not the Aurora cluster configuration used by option C's database activity stream. Therefore, it is not the correct pairing for the required audit solution.

https://docs.aws.amazon.com/en_en/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.Encryption.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.SSL.html

✓ **C** Create a database activity stream in the database cluster. Use Amazon Data Firehose to send the data to an Amazon S3 bucket. Enable S3 Object Lock in compliance mode with a 10-year retention period.
Aurora database activity streams provide activity records that can be delivered through Firehose to S3, outside the control of database users. S3 Object Lock in compliance mode protects the audit records from being deleted or overwritten during the specified 10-year retention period.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html
https://docs.aws.amazon.com/us_en/AmazonS3/latest/userguide/object-lock-managing.html
https://docs.aws.amazon.com/en_en/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.Monitoring.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.AuditLog.html
https://docs.aws.amazon.com/streams/latest/dev/using-other-services-aurora.html

✗ **D** Install pgAudit in the database cluster. Store the data in Amazon CloudWatch Logs. Use Amazon Data Firehose to send the data to an Amazon S3 bucket. Enable S3 Object Lock in compliance mode with a 10-year retention period.
The compliance-mode S3 destination would protect records after they are written there, but pgAudit is not the database activity stream specified for the Aurora cluster solution. It does not provide the required managed activity-stream audit path for this combination.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html
https://docs.aws.amazon.com/us_en/AmazonS3/latest/userguide/object-lock-managing.html
https://docs.aws.amazon.com/en_en/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.AuditLog.html

✗ **E** Install pgAudit in the database cluster. Store the data in Amazon CloudWatch Logs. Use Amazon Data Firehose to send the data to an Amazon S3 bucket. Enable S3 Object Lock in governance mode with a 10-year retention period.
Governance mode does not provide the required level of immutability: users with appropriate permissions can bypass retention protections. The requirement calls for audit data that cannot be tampered with, so compliance mode is needed; this option also uses pgAudit rather than the Aurora activity-stream approach.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html
https://docs.aws.amazon.com/us_en/AmazonS3/latest/userguide/object-lock-managing.html
https://docs.aws.amazon.com/en_en/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.html
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.AuditLog.html

Question 8

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A company is deploying an AI agent on Amazon Bedrock AgentCore Runtime. The agent automates business processes by accessing AWS resources and external services. The agent performs actions that do not require user interaction.

The company needs the agent to authenticate as its own identity when accessing downstream services. The agent must maintain an audit trail that logs the agent identity for every action the agent performs.

Which solution will meet these requirements?

A. Register the agent as a workload identity in an Amazon Bedrock AgentCore Identity agent identity directory. Configure resource credential providers by using the OAuth 2.0 client credentials grant flow for each downstream service. Use the workload identity to give downstream services and audit systems the ability to identify which agent performed each action.

B. Create an IAM user for the agent. Store long-lived access keys in AWS Secrets Manager. Configure the agent to retrieve the keys at runtime and use the keys to authenticate to each downstream service.

C. Register the agent as a workload identity in an Amazon Bedrock AgentCore Identity agent identity directory. Configure resource credential providers by using the OAuth 2.0 authorization code grant flow for each downstream service to obtain access tokens delegated on behalf of each user.

D. Create an IAM role for the agent. Configure each downstream service to accept IAM Signature Version 4 (SigV4) signed requests. Use AWS CloudTrail to log the IAM role ARN for auditing purposes.

Answer: A. Register the agent as a workload identity in an Amazon Bedrock AgentCore Identity agent identity directory. Configure resource credential providers by using the OAuth 2.0 client credentials grant flow for each downstream service. Use the workload identity to give downstream services and audit systems the ability to identify which agent performed each action.

Explanation

The agent needs a noninteractive identity of its own that downstream services can use to identify it and associate actions with it. Registering it as a workload identity and configuring OAuth 2.0 client credentials providers supports machine-to-machine access without user delegation. This lets the agent authenticate as itself and supports auditing actions by agent identity.

https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/identity.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/identity-add-oauth-client-custom.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/understanding-agent-identities.html

✓ **A** Register the agent as a workload identity in an Amazon Bedrock AgentCore Identity agent identity directory. Configure resource credential providers by using the OAuth 2.0 client credentials grant flow for each downstream service. Use the workload identity to give downstream services and audit systems the ability to identify which agent performed each action.
This fits an agent acting autonomously: the client credentials grant is for noninteractive, machine-to-machine authentication rather than a user signing in. The workload identity gives the agent its own identity, which can be used to associate its downstream actions with that agent for auditing.

https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/identity.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/identity-add-oauth-client-custom.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/understanding-agent-identities.html

✗ **B** Create an IAM user for the agent. Store long-lived access keys in AWS Secrets Manager. Configure the

Free demo · full exam: 681 questions

agent to retrieve the keys at runtime and use the keys to authenticate to each downstream service. This uses long-lived IAM user credentials rather than an AgentCore workload identity, so downstream systems and audit records would not identify the agent through its own workload identity. It also does not provide a suitable general authentication method for external services that do not accept those AWS credentials.

https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/identity.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/understanding-agent-identities.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/security-iam.html

✗ **C** Register the agent as a workload identity in an Amazon Bedrock AgentCore Identity agent identity directory. Configure resource credential providers by using the OAuth 2.0 authorization code grant flow for each downstream service to obtain access tokens delegated on behalf of each user.
The authorization code flow obtains delegated access on behalf of a user and involves user authorization. The scenario explicitly requires actions that do not require user interaction and requires the agent to authenticate as its own identity, so client credentials is the better fit.

https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/identity-add-oauth-client-custom.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/common-use-cases.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/oauth2-authorization-url-session-binding.html

✗ **D** Create an IAM role for the agent. Configure each downstream service to accept IAM Signature Version 4 (SigV4) signed requests. Use AWS CloudTrail to log the IAM role ARN for auditing purposes.
An IAM role ARN identifies the AWS role used for signed requests, not necessarily the individual AgentCore agent identity. This approach also depends on downstream services accepting SigV4, and CloudTrail role logging does not by itself establish the required agent identity for every action across AWS resources and external services.

https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/understanding-agent-identities.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/security-iam.html

Question 9

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Determine appropriate data and network protection controls.

A global company operates a platform that serves customers across multiple AWS Regions. The platform stores customer behavioral data.

For data residency compliance, the company must ensure that personally identifiable information (PII) data remains within the Region where the data is collected. Additionally, the company must ensure that cross-Region data analysis uses only anonymized datasets.

Which solution will meet these requirements?

A. Deploy AWS Outposts in each Region to keep data on premises. Store data in Amazon S3 on Outposts. Use AWS Glue DataBrew to anonymize PII data. Analyze cross-Region data by using Amazon Athena.

B. Deploy Amazon Aurora PostgreSQL clusters in separate Regions. Use AWS Glue DataBrew to anonymize PII data. Analyze cross-Region data by using Amazon Redshift Serverless.

C. Deploy Amazon Aurora PostgreSQL clusters in separate Regions. Use AWS Lambda functions to anonymize PII data before replication. Use AWS PrivateLink to connect Amazon QuickSight to cross-Region databases for analysis.

D. Deploy Amazon S3 buckets in each Region. Enable S3 Block Public Access and bucket policies to prevent cross-Region replication. Use Amazon Macie to anonymize data. Analyze cross-Region data by using Amazon Athena.

Free demo · full exam: 681 questions

Answer: B. Deploy Amazon Aurora PostgreSQL clusters in separate Regions. Use AWS Glue DataBrew to anonymize PII data. Analyze cross-Region data by using Amazon Redshift Serverless.

Explanation

B keeps customer data in separate Aurora PostgreSQL clusters in the Regions where it is collected, while AWS Glue DataBrew anonymizes PII before datasets are used for cross-Region analysis. Amazon Redshift Serverless supports cross-Region data sharing, providing an analytics path for the anonymized datasets. This addresses both regional data residency and the requirement that cross-Region analysis use only anonymized data.

https://docs.aws.amazon.com/databrew/latest/dg/personal-information-protection.html
https://docs.aws.amazon.com/en_en/AmazonRDS/latest/AuroraUserGuide/Concepts.RegionsAndAvailabilityZones.html
https://docs.aws.amazon.com/redshift/latest/mgmt/serverless-datasharing.html

✗ **A** Deploy AWS Outposts in each Region to keep data on premises. Store data in Amazon S3 on Outposts. Use AWS Glue DataBrew to anonymize PII data. Analyze cross-Region data by using Amazon Athena.
DataBrew can anonymize PII, but this option does not specify a cross-Region sharing arrangement that makes only the anonymized datasets available for analysis. Deploying Outposts and storing data on premises does not itself provide that controlled cross-Region analytics path; Redshift Serverless data sharing is the fit for that part of the requirement.

https://docs.aws.amazon.com/databrew/latest/dg/personal-information-protection.html
https://docs.aws.amazon.com/redshift/latest/mgmt/serverless-datasharing.html

✓ **B** Deploy Amazon Aurora PostgreSQL clusters in separate Regions. Use AWS Glue DataBrew to anonymize PII data. Analyze cross-Region data by using Amazon Redshift Serverless.
Separate Aurora clusters let the company keep collected customer data in its Region. DataBrew can anonymize PII, and Redshift Serverless supports cross-Region data sharing so the analysis can use anonymized datasets rather than regional PII.

https://docs.aws.amazon.com/databrew/latest/dg/personal-information-protection.html
https://docs.aws.amazon.com/en_en/AmazonRDS/latest/AuroraUserGuide/Concepts.RegionsAndAvailabilityZones.html
https://docs.aws.amazon.com/redshift/latest/mgmt/serverless-datasharing.html

✗ **C** Deploy Amazon Aurora PostgreSQL clusters in separate Regions. Use AWS Lambda functions to anonymize PII data before replication. Use AWS PrivateLink to connect Amazon QuickSight to cross-Region databases for analysis.
The separate regional clusters and anonymization before replication address residency and reduce the risk of replicating PII. However, PrivateLink provides private connectivity; it is not itself a cross-Region analytics or data-sharing service for QuickSight. Redshift Serverless provides a supported cross-Region data-sharing approach for the analysis.

https://docs.aws.amazon.com/en_en/AmazonRDS/latest/AuroraUserGuide/Concepts.RegionsAndAvailabilityZones.html
https://docs.aws.amazon.com/redshift/latest/mgmt/serverless-datasharing.html

✗ **D** Deploy Amazon S3 buckets in each Region. Enable S3 Block Public Access and bucket policies to prevent cross-Region replication. Use Amazon Macie to anonymize data. Analyze cross-Region data by using Amazon Athena.
Preventing cross-Region replication leaves the data in separate Regions without a mechanism in this option to make datasets available for cross-Region analysis. Also, Macie is used to discover and classify sensitive data, not to anonymize it; therefore it does not meet the stated anonymization requirement.

Free demo · full exam: 681 questions

Question 10

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design solutions that provide access to AWS accounts.

A company is creating a solution that can move 400 employees into a remote working environment in the event of an unexpected disaster. The user desktops have a mix of Windows and Linux operating systems. Multiple types of software, such as web browsers and mail clients, are installed on each desktop.

A solutions architect needs to implement a solution that can be integrated with the company's onpremises Active Directory to allow employees to use their existing identity credentials. The solution must provide multi-factor authentication (MFA) and must replicate the user experience from the existing desktops.

Which solution will meet these requirements?

A. Use Amazon WorkSpaces for the cloud desktop service. Set up a VPN connection to the on-premises network. Create an AD Connector, and connect to the on-premises Active Directory. Activate MFA for Amazon WorkSpaces by using the AWS Management Console.

B. Use Amazon AppStream 2.0 as an application streaming service. Configure Desktop View for the employees. Set up a VPN connection to the on-premises network. Set up Active Directory Federation Services (AD FS) on premises. Connect the VPC network to AD FS through the VPN connection.

C. Use Amazon WorkSpaces for the cloud desktop service. Set up a VPN connection to the on-premises network. Create an AD Connector, and connect to the on-premises Active Directory. Configure a RADIUS server for MFA.

D. Use Amazon AppStream 2.0 as an application streaming service. Set up Active Directory Federation Services on premises. Configure MFA to grant users access on AppStream 2.0.

Answer: C. Use Amazon WorkSpaces for the cloud desktop service. Set up a VPN connection to the on- premises network. Create an AD Connector, and connect to the on-premises Active Directory. Configure a RADIUS server for MFA.

Explanation

Amazon WorkSpaces provides cloud desktops, which fits the requirement to reproduce employees' desktop environment rather than merely stream selected applications. An AD Connector lets WorkSpaces use the company's on-premises Active Directory, so employees can use their existing credentials. To add MFA to WorkSpaces, configure a RADIUS server; this makes option C the solution that meets the requirements.

https://docs.aws.amazon.com/workspaces/latest/adminguide/connect-mfa.html
https://docs.aws.amazon.com/workspaces/latest/adminguide/
https://docs.aws.amazon.com/whitepapers/latest/best-practices-deploying-amazon-workspaces/ad-connector-role-with-workspaces.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **A** Use Amazon WorkSpaces for the cloud desktop service. Set up a VPN connection to the on-premises network. Create an AD Connector, and connect to the on-premises Active Directory. Activate MFA for Amazon WorkSpaces by using the AWS Management Console.
WorkSpaces with an AD Connector fits the cloud desktop and existing-credentials requirements. However, WorkSpaces MFA is configured using a RADIUS server, not by simply activating MFA in the AWS Management Console, so this option does not specify the required MFA setup.

https://docs.aws.amazon.com/workspaces/latest/adminguide/connect-mfa.html
https://docs.aws.amazon.com/whitepapers/latest/best-practices-deploying-amazon-workspaces/ad-connector-role-with-workspaces.html?trk=article-ssr-frontend-pulse_little-text-block

Free demo · full exam: 681 questions

✗ **B** Use Amazon AppStream 2.0 as an application streaming service. Configure Desktop View for the employees. Set up a VPN connection to the on-premises network. Set up Active Directory Federation Services (AD FS) on premises. Connect the VPC network to AD FS through the VPN connection.
AppStream 2.0 is presented here as an application-streaming solution, whereas the requirement is for a cloud desktop that reproduces the existing desktop environment. This option also describes connecting to AD FS but does not specify configuring MFA for access.

✓ **C** Use Amazon WorkSpaces for the cloud desktop service. Set up a VPN connection to the on-premises network. Create an AD Connector, and connect to the on-premises Active Directory. Configure a RADIUS server for MFA.
WorkSpaces provides the cloud desktop experience needed to make employees' existing desktop environment available remotely. The AD Connector connects WorkSpaces to the company's on-premises Active Directory for existing identity credentials, and a RADIUS server provides the required MFA.

https://docs.aws.amazon.com/workspaces/latest/adminguide/connect-mfa.html
https://docs.aws.amazon.com/workspaces/latest/adminguide/
https://docs.aws.amazon.com/whitepapers/latest/best-practices-deploying-amazon-workspaces/ad-connector-role-with-workspaces.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **D** Use Amazon AppStream 2.0 as an application streaming service. Set up Active Directory Federation Services on premises. Configure MFA to grant users access on AppStream 2.0.
Although the option specifies MFA through AD FS, AppStream 2.0 is an application-streaming service rather than the cloud desktop solution called for to reproduce the existing desktop environment. It also does not describe connecting the solution to the on-premises Active Directory as the WorkSpaces/AD Connector approach does.

Dumpsgate

Question 11

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A company uses a single organization in AWS Organizations with all features enabled to manage 300 AWS accounts. A solutions architect must ensure that all private Amazon ECR repositories that are deployed in the member accounts are accessible only to principals that belong to the same organization. The company must enforce access controls across the entire organization and for all current and future ECR repositories.

Which solution will meet these requirements?

A. Attach an SCP to the organization root that denies all ecr:* actions unless the aws:PrincipalType property has a value of Service.

B. Attach a resource control policy (RCP) to the organization root that allows access to ECR only when the aws:PrincipalOrgID condition matches the company's organization ID.

C. Modify IAM policies across all accounts to include a deny statement for ecr:* unless the caller's aws:PrincipalAccount property matches the organization's ID.

D. Share all ECR repositories through AWS Resource Access Manager (AWS RAM). Control principal access based on account-level allowlists.

Answer: B. Attach a resource control policy (RCP) to the organization root that allows access to ECR only when the aws:PrincipalOrgID condition matches the company's organization ID.

Explanation

An RCP attached to the organization root provides an organization-wide resource access control for supported services, including ECR. Its condition can restrict access to principals whose `aws:PrincipalOrgID` matches the company's organization ID, applying across member accounts and their ECR repositories, including repositories created in the future. The RCP limits access; it does not replace the IAM permissions otherwise needed for a request to succeed.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_rcps.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_authorization_policies.html

✗ **A** Attach an SCP to the organization root that denies all ecr:* actions unless the aws:PrincipalType property has a value of Service.
An SCP restricts permissions for principals in the organization; it does not by itself control access by principals outside the organization to resources in member accounts. Also, allowing only service principals would not express the requirement to allow principals from the same organization and could block ordinary account principals.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_policies_create.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_authorization_policies.html

✓ **B** Attach a resource control policy (RCP) to the organization root that allows access to ECR only when the aws:PrincipalOrgID condition matches the company's organization ID.
RCPs let an organization define resource-side access controls for supported services across accounts in the organization. The `aws:PrincipalOrgID` condition restricts ECR access to principals belonging to the specified organization, and attaching the policy at the root covers current and future member accounts and their repositories.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_rcps.html
https://docs.aws.amazon.com/AmazonECR/latest/userguide/registry-permissions-examples.html?icmpid=docs_ecr_hp-registry-private

✗ **C** Modify IAM policies across all accounts to include a deny statement for ecr:* unless the caller's aws:PrincipalAccount property matches the organization's ID.
IAM identity policies in member accounts do not provide a central resource-side restriction against principals outside the organization. Further, `aws:PrincipalAccount` identifies an AWS account, whereas an organization ID is a different identifier, so comparing the account value to the organization ID does not implement the intended test.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_authorization_policies.html
https://docs.aws.amazon.com/AmazonECR/latest/userguide/security_iam_id-based-policy-examples.html
https://docs.aws.amazon.com/service-authorization/latest/reference/list_ecr.html

✗ **D** Share all ECR repositories through AWS Resource Access Manager (AWS RAM). Control principal access based on account-level allowlists.
This approach would require managing repository sharing and account allowlists rather than enforcing the organization-wide condition for all current and future repositories. ECR repository access is governed through ECR permissions, and an organization-root RCP is the centralized control suited to this requirement.

https://docs.aws.amazon.com/AmazonECR/latest/userguide/repository-policies.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_rcps.html

Free demo · full exam: 681 questions

Question 12

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain thirdparty software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement . The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers Other IAM users groups, roles, and account administrators in the company should be denied Private Marketplace administrative access What is the MOST efficient way to design an architecture to meet these requirements?

A. Create an IAM role named procurement-manager-role in all AWS accounts in the organization Add the PowerUserAccess managed policy to the role Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPrivateMarketplaceAdminFullAccess managed policy.

B. Create an IAM role named procurement-manager-role in all AWS accounts in the organization Add the AdministratorAccess managed policy to the role Define a permissions boundary with the AWSPrivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.

C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.

D. Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developers. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an SCP inOrganizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Apply the SCP to all the shared services accounts in the organization.

Answer: C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.

Explanation

Option C places the procurement role in the shared services accounts used by procurement managers and grants it the AWS-managed policy for Private Marketplace administration. Organization-level SCPs can restrict that administrative access across the organization to the designated role, while separately preventing others from creating a role with the trusted name. This fits the requirement to let procurement managers manage the approved list while denying administrative access to other users, roles, and account administrators. [1][3][9][12]

https://docs.aws.amazon.com/aws-managed-policy/latest/reference/AWSPrivateMarketplaceAdminFullAccess.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-private-marketplace.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_policies_create.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_authorization_policies.html

✗ **A** Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the PowerUserAccess managed policy to the role. Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPrivateMarketplaceAdminFullAccess managed policy.
PowerUserAccess is not the AWS-managed policy that grants Private Marketplace administration, so this role would not provide the specified administrative access. A deny attached individually to users and roles is also less efficient and does not establish the organization-wide guardrail needed to restrict account administrators; IAM policies control actions, not permissions to use a managed policy as though it were an action.

https://docs.aws.amazon.com/aws-managed-policy/latest/reference/AWSPrivateMarketplaceAdminFullAccess.html
https://docs.aws.amazon.com/marketplace/latest/buyerguide/buyer-security-iam-awsmanpol.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_policies_create.html

✗ **B** Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the AdministratorAccess managed policy to the role. Define a permissions boundary with the AWSPrivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.
AdministratorAccess gives the procurement role far broader permissions than needed, and a permissions boundary on developer roles does not by itself grant developers Private Marketplace permissions or deny administrators those permissions. The design also fails to use an organization-wide SCP to enforce the requested restriction across accounts.

https://docs.aws.amazon.com/aws-managed-policy/latest/reference/AWSPrivateMarketplaceAdminFullAccess.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_policies_create.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_authorization_policies.html

✓ **C** Create an IAM role named procurement-manager-role in all the shared services accounts in the organization. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.
The designated role is placed in the shared services accounts used by procurement managers and receives the managed policy intended for Private Marketplace administration. The root-level SCPs enforce the restriction across the organization and prevent other principals from creating a role with the trusted name, addressing the risk of bypassing the role-based restriction. SCPs set organization-level permission guardrails; the role's IAM policy supplies its allowed administrative permissions.

https://docs.aws.amazon.com/aws-managed-policy/latest/reference/AWSPrivateMarketplaceAdminFullAccess.html
https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-private-marketplace.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_policies_create.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_authorization_policies.html

✗ **D** Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developers. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an SCP in Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Apply the SCP to all the shared services accounts in the organization.
The role is created in developer accounts rather than the shared services accounts intended for procurement managers. More importantly, attaching the SCP only to shared services accounts leaves accounts outside that scope without the organization-wide restriction, so it does not ensure that other account administrators and principals are denied administrative access throughout the organization.

https://docs.aws.amazon.com/aws-managed-policy/latest/reference/AWSPrivateMarketplaceAdminFullAccess.html
https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-private-marketplace.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_policies_create.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_authorization_policies.html

Question 13

Multiple choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A research center is migrating to the AWS Cloud and has moved its on-premises 1 PB object storage to an Amazon S3 bucket. One hundred scientists are using this object storage to store their workrelated documents. Each scientist has a personal folder on the object store. All the scientists are members of a single IAMuser group.

The research center's compliance officer is worried that scientists will be able to access each other's work. The research center has a strict obligation to report on which scientist accesses which documents.

The team that is responsible for these reports has little AWS experience and wants a ready-to-use solution that minimizes operational overhead.

Which combination of actions should a solutions architect take to meet these requirements? (Select TWO.)

A. Create an identity policy that grants the user read and write access. Add a condition that specifies that the S3 paths must be prefixed with ${aws:username}. Apply the policy on the scientists' IAM user group.

B. Configure a trail with AWS CloudTrail to capture all object-level events in the S3 bucket. Store the trail output in another S3 bucket. Use Amazon Athena to query the logs and generate reports.

C. Enable S3 server access logging. Configure another S3 bucket as the target for log delivery. Use Amazon Athena to query the logs and generate reports.

D. Create an S3 bucket policy that grants read and write access to users in the scientists' IAM user group.

E. Configure a trail with AWS CloudTrail to capture all object-level events in the S3 bucket and write the events to Amazon CloudWatch. Use the Amazon Athena CloudWatch connector to query the logs and generate reports.

Answer

- A. **Create an identity policy that grants the user read and write access. Add a condition that specifies that the S3 paths must be prefixed with ${aws:username}. Apply the policy on the scientists' IAM user group.**
- B. **Configure a trail with AWS CloudTrail to capture all object-level events in the S3 bucket. Store the trail output in another S3 bucket. Use Amazon Athena to query the logs and generate reports.**

Explanation

Apply a group identity policy that scopes each scientist’s S3 access to a prefix based on that scientist’s IAM username, so scientists can access their own folders rather than one another’s. For compliance reporting, use CloudTrail object-level data events and store the trail in S3, where the reporting team can query the events with Athena. This provides an auditable record of object access and a straightforward path to reports without adding a CloudWatch-to-Athena connector.

https://docs.aws.amazon.com/athena/latest/ug/query-examples-cloudtrail-logs.html
https://docs.aws.amazon.com/athena/latest/ug/monitor-with-cloudtrail.html
https://docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_s3_home-directory-console.html

✓ **A** Create an identity policy that grants the user read and write access. Add a condition that specifies that the S3 paths must be prefixed with ${aws:username}. Apply the policy on the scientists' IAM user

Free demo · full exam: 681 questions

group.
The username-based prefix scopes each scientist's access to their own folder, addressing the concern that they could access each other's work. Applying the policy to the shared group lets the scientists use the same group while the policy evaluates the individual user's username. AWS documents this kind of S3 home-directory policy pattern.

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_s3_home-directory-console.html

✓ **B** Configure a trail with AWS CloudTrail to capture all object-level events in the S3 bucket. Store the trail output in another S3 bucket. Use Amazon Athena to query the logs and generate reports.
CloudTrail object-level data events record activity on S3 objects, providing the per-access information needed to report which scientist accessed which documents. Storing the trail in S3 allows Athena to query CloudTrail logs and generate reports without requiring a separate log-querying connector.

https://docs.aws.amazon.com/athena/latest/ug/query-examples-cloudtrail-logs.html
https://docs.aws.amazon.com/athena/latest/ug/monitor-with-cloudtrail.html
https://docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html

✗ **C** Enable S3 server access logging. Configure another S3 bucket as the target for log delivery. Use Amazon Athena to query the logs and generate reports.
S3 server access logs can be delivered to a bucket and queried, but this is not the selected approach for the requested audit reports. CloudTrail object-level events paired with Athena provide the specified, direct route to querying records of object access by the requesting identity.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/using-s3-access-logs-to-identify-requests.html
https://docs.aws.amazon.com/us_en/athena/latest/ug/querying-apache-logs.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerLogs.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/LogFormat.html

✗ **D** Create an S3 bucket policy that grants read and write access to users in the scientists' IAM user group.
Granting the whole group read and write access to the bucket does not limit each scientist to their personal folder. As written, it does not meet the requirement to prevent access to other scientists' work.

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_s3_home-directory-console.html

✗ **E** Configure a trail with AWS CloudTrail to capture all object-level events in the S3 bucket and write the events to Amazon CloudWatch. Use the Amazon Athena CloudWatch connector to query the logs and generate reports.
CloudTrail can be queried through Athena, but this option adds a CloudWatch Logs destination and an Athena connector rather than using the simpler CloudTrail-to-S3 and Athena approach. The additional integration is unnecessary for the team's goal of minimizing operational overhead.

https://docs.aws.amazon.com/athena/latest/ug/connectors-cloudwatch.html
https://docs.aws.amazon.com/athena/latest/ug/monitor-with-cloudtrail.html
https://docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html

Free demo · full exam: 681 questions

Question 14

Multiple choice · Domain: D1 Design Solutions for Organizational Complexity › Design solutions that provide access to AWS accounts.

A publishing company's design team updates the icons and other static assets that an ecommerce web application uses. The company serves the icons and assets from an Amazon S3 bucket that is hosted in the company's production account. The company also uses a development account that members of the design team canaccess.

After the design team tests the static assets in the development account, the design team needs to load the assets into the S3 bucket in the production account. A solutions architect must provide the design team with access to the production account without exposing other parts of the web application to the risk of unwanted changes.

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

A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket.

B. In the development account, create a new IAM policy that allows read and write access to the S3 bucket.

C. In the production account, create a role. Attach the new policy to the role. Define the development account as a trusted entity.

D. In the development account, create a role. Attach the new policy to the role. Define the production account as a trusted entity.

E. In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account.

F. In the development account, create a group that contains all tfje IAM users of the design team. Attach a different IAM policy to the group to allow the sts;AssumeRole action on the role in the development account.

Answer

- A. **In the production account, create a new IAM policy that allows read and write access to the S3 bucket.**
- C. **In the production account, create a role. Attach the new policy to the role. Define the development account as a trusted entity.**
- E. **In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account.**

Explanation

Create the S3 read/write policy in the production account and attach it to a role that trusts the development account. Then grant the design team’s development-account users permission to assume that production role. This provides access to the production S3 assets through a specific role rather than granting access to other parts of the web application.

https://docs.aws.amazon.com/us_en/IAM/latest/UserGuide/access_policies-cross-account-resource-access.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic-cross-account.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_update-role-trust-policy.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_permissions-to-switch.html

✓ **A** In the production account, create a new IAM policy that allows read and write access to the S3

Free demo · full exam: 681 questions

bucket.
The role in the production account needs permissions to read and write the production S3 bucket. Scope the policy to that bucket so the design team can update the assets without gaining unrelated application permissions.

https://docs.aws.amazon.com/us_en/IAM/latest/UserGuide/access_policies-cross-account-resource-access.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic-cross-account.html
https://docs.aws.amazon.com/us_en/IAM/latest/UserGuide/access_controlling.html

✗ **B** In the development account, create a new IAM policy that allows read and write access to the S3 bucket.
A policy in the development account alone does not grant permissions to the production-account role to access the bucket. The bucket-access policy belongs in the production account and should be attached to the role that the team assumes.

https://docs.aws.amazon.com/us_en/IAM/latest/UserGuide/access_policies-cross-account-resource-access.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic-cross-account.html

✓ **C** In the production account, create a role. Attach the new policy to the role. Define the development account as a trusted entity.
The production role can carry the S3 permissions from option A and trust the development account to assume it. This is the cross-account role setup that gives the team the required access while keeping permissions limited to the role's policy.

https://docs.aws.amazon.com/us_en/IAM/latest/UserGuide/access_policies-cross-account-resource-access.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_update-role-trust-policy.html

✗ **D** In the development account, create a role. Attach the new policy to the role. Define the production account as a trusted entity.
This puts the role in the wrong account and makes the production account trusted by the development account's role. It does not provide the intended role in production with permissions to access the production S3 bucket.

https://docs.aws.amazon.com/us_en/IAM/latest/UserGuide/access_policies-cross-account-resource-access.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_update-role-trust-policy.html

✓ **E** In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account.
The development-account users need identity-based permission to call sts:AssumeRole on the production role. Applying that permission to a group for the design team grants the needed access to its members without requiring the permission for unrelated users.

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic-cross-account.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_permissions-to-switch.html

✗ **F** In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the development account.
The group's permission needs to target the role in the production account, not a role in the development account. Assuming a development-account role would not provide the required cross-account access to the production S3 bucket.

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic-cross-account.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_permissions-to-switch.html

Free demo · full exam: 681 questions

Question 15

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design solutions that provide access to AWS accounts.

A company needs its remote employees to access the web interface of an application. The company hosts the application in a VPC in the AWS Cloud. The application runs on a fleet of Amazon EC2 instances in private subnets across multiple Availability Zones.

The remote employees are not currently using a corporate VPN connection. Company policy does not allow a direct connection from the public internet to the application.

Which solution will meet these requirements?

A. Create a customer gateway by using the company's external IP address. Create a virtual private gateway, and attach it to the VPC. Create an AWS Site-to-Site VPN connection between the customer gateway and the virtual private gateway. Modify the route table to route corporate traffic back to the customer gateway.

B. Generate server and client certificates and keys. Create an AWS Client VPN endpoint. Associate the private subnets with the Client VPN endpoint. Authorize the remote employees to access the VPC. Download the Client VPN endpoint configuration file to use in the remote employees' VPN client application.

C. Create a transit gateway in the VPC. Associate the private subnets with the transit gateway. Enable AWS Client VPN connection endpoints on the transit gateway. Deploy the Client VPN software to each of the remote employees. Configure route tables to route traffic between the Client VPN connections and the private subnets.

D. Create a network ACL. Add inbound and outbound rules for the IP addresses of the remote employees. Associate the private subnets with the network ACL. Update the existing security group that is associated with the fleet of EC2 instances to permit the IP addresses of the remote employees.

Answer: B. Generate server and client certificates and keys. Create an AWS Client VPN endpoint. Associate the private subnets with the Client VPN endpoint. Authorize the remote employees to access the VPC. Download the Client VPN endpoint configuration file to use in the remote employees' VPN client application.

Explanation

AWS Client VPN gives remote employees a VPN connection into the VPC without requiring a corporate Site-to-Site VPN. Associating the Client VPN endpoint with private subnets lets connected employees reach the application instances there, while the application need not accept direct connections from the public internet. The endpoint configuration file provides the settings employees need for their VPN client, and the endpoint can be authorized to allow their access.

https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/cvpn-working-target-associate.html

✗ **A** Create a customer gateway by using the company's external IP address. Create a virtual private gateway, and attach it to the VPC. Create an AWS Site-to-Site VPN connection between the customer gateway and the virtual private gateway. Modify the route table to route corporate traffic back to the customer gateway.
A Site-to-Site VPN connects a corporate network to an AWS VPC through customer-gateway and virtual-private-gateway devices. The scenario says remote employees are not using a corporate VPN, so this does not provide the requested individual remote-user access without setting up that corporate network connection.

https://docs.aws.amazon.com/vpn/latest/s2svpn/how_it_works.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html

✓ **B** Generate server and client certificates and keys. Create an AWS Client VPN endpoint. Associate the private subnets with the Client VPN endpoint. Authorize the remote employees to access the VPC. Download the Client VPN endpoint configuration file to use in the remote employees' VPN client application.
AWS Client VPN is designed to let remote users establish VPN connections to resources in a VPC. Associating private subnets with the endpoint provides a network path to the application instances, and authorization rules control which users can reach the VPC. Employees can use the downloaded endpoint configuration in their VPN client.

https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/cvpn-working-target-associate.html

✗ **C** Create a transit gateway in the VPC. Associate the private subnets with the transit gateway. Enable AWS Client VPN connection endpoints on the transit gateway. Deploy the Client VPN software to each of the remote employees. Configure route tables to route traffic between the Client VPN connections and the private subnets.
A transit gateway is not a feature that enables Client VPN endpoints on itself. Client VPN endpoints are associated with target subnets; a transit gateway may be used in certain network designs to route traffic, but it is not required for the direct VPC access described here.

https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/cvpn-working-target-associate.html
https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/cvpn-tgw.html
https://docs.aws.amazon.com/vpc/latest/tgw/tgw-transit-gateways.html

✗ **D** Create a network ACL. Add inbound and outbound rules for the IP addresses of the remote employees. Associate the private subnets with the network ACL. Update the existing security group that is associated with the fleet of EC2 instances to permit the IP addresses of the remote employees.
Network ACL and security-group rules filter traffic; they do not create a secure remote-access connection into the VPC. Allowing employee public IP addresses would still depend on a direct internet path, which company policy disallows, and the instances are in private subnets.

Dumpsgate

Question 16

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Determine appropriate data and network protection controls.

A company deploys its applications on Amazon EC2 instances in a VPC in the us-east-1 Region. The company recently acquired another company that has a data center in the ap-east-1 Region.

A private application in us-east-1 needs to access the data that is stored in the data center in ap-east1. The company needs a solution that supports IPv6. The solution also must encrypt the data in transit.

The company has already created an AWS customer gateway that represents the customer gateway device in the data center in ap-east-1. The company will configure the route tables and modify the existing security groups to allow appropriate traffic after the network components are created.

Which solution will meet the requirements for connectivity?

A. Create a transit gateway. Attach the transit gateway to the VPC in us-east-1. Create an AWS Site-to-Site VPN connection. Specify the transit gateway as the target gateway. Download the configuration file and set up the customer gateway device.

B. Create a virtual private gateway. Attach the virtual private gateway to the VPC in us-east-1. Create an AWS Client VPN endpoint. Configure the customer gateway device to connect to the Client VPN endpoint.

C. Create a virtual private gateway. Attach the virtual private gateway to the VPC in us-east-1. Create an AWS Site-to-Site VPN connection. Specify the virtual private gateway as the target gateway. Download the configuration file and set up the customer gateway device.

D. Create a transit gateway. Attach the transit gateway to the VPC in us-east-1. Create an AWS Client VPN endpoint. Configure the customer gateway device to connect to the Client VPN endpoint.

Answer: A. Create a transit gateway. Attach the transit gateway to the VPC in us-east-1. Create an AWS Site-to- Site VPN connection. Specify the transit gateway as the target gateway. Download the configuration file and set up the customer gateway device.

Explanation

A transit gateway can attach to the VPC in us-east-1 and serve as the target for an AWS Site-to-Site VPN connection to the data center's customer gateway. Site-to-Site VPN provides encrypted traffic in transit, and the IPv6 requirements are met by using the transit-gateway VPN configuration that supports IPv6. The downloaded configuration file provides the settings needed to configure the customer gateway device.

https://docs.aws.amazon.com/vpn/latest/s2svpn/ipv4-ipv6.html
https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpc-attachments.html
https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html
https://docs.aws.amazon.com/vpn/latest/s2svpn/data-protection.html

✓ **A** Create a transit gateway. Attach the transit gateway to the VPC in us-east-1. Create an AWS Site-to-Site VPN connection. Specify the transit gateway as the target gateway. Download the configuration file and set up the customer gateway device.
This creates a network-to-network Site-to-Site VPN between the VPC and the data center's customer gateway, with the transit gateway attached to the VPC. The transit-gateway VPN option supports IPv6, and the VPN encrypts traffic in transit; the downloaded configuration helps set up the customer gateway device.

https://docs.aws.amazon.com/vpn/latest/s2svpn/ipv4-ipv6.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpc-attachments.html
https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html
https://docs.aws.amazon.com/vpn/latest/s2svpn/data-protection.html

✗ **B** Create a virtual private gateway. Attach the virtual private gateway to the VPC in us-east-1. Create an AWS Client VPN endpoint. Configure the customer gateway device to connect to the Client VPN endpoint.
Client VPN is for client access to AWS resources, not for establishing the required site-to-site connection between a VPC and a customer gateway device at a data center. This option therefore does not provide the appropriate connectivity for the scenario.

https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/cvpn-working-endpoints.html

✗ **C** Create a virtual private gateway. Attach the virtual private gateway to the VPC in us-east-1. Create an AWS Site-to-Site VPN connection. Specify the virtual private gateway as the target gateway. Download the configuration file and set up the customer gateway device.
A Site-to-Site VPN is the appropriate type of connection, but this option targets a virtual private gateway. The IPv6 Site-to-Site VPN capability required here is supported with a transit gateway, so this choice does not meet the IPv6 requirement.

https://docs.aws.amazon.com/vpn/latest/s2svpn/ipv4-ipv6.html
https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html

✗ **D** Create a transit gateway. Attach the transit gateway to the VPC in us-east-1. Create an AWS Client VPN endpoint. Configure the customer gateway device to connect to the Client VPN endpoint.
Attaching a transit gateway to the VPC does not make a Client VPN endpoint a site-to-site connection. Client VPN is intended for client access, so this option does not connect the data center's customer gateway device to the VPC as required.

https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpc-attachments.html
https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/cvpn-working-endpoints.html

Dumpsgate

Question 17

Multiple choice · Domain: D1 Design Solutions for Organizational Complexity › Design solutions that provide access to AWS accounts.

A company has many AWS accounts and uses AWS Organizations to manage all of them. A solutions architect must implement a solution that the company can use to share a common network across multiple accounts.

The company's infrastructure team has a dedicated infrastructure account that has a VPC. The infrastructure team must use this account to manage the network. Individual accounts cannot have the ability to manage their own networks. However, individual accounts must be able to create AWS resources within subnets.

Which combination of actions should the solutions architect perform to meet these requirements? (Select TWO.)

A. Create a transit gateway in the infrastructure account.

B. Enable resource sharing from the AWS Organizations management account.

C. Create VPCs in each AWS account within the organization in AWS Organizations. Configure the VPCs to share the same CIDR range and subnets as the VPC in the infrastructure account. Peer the VPCs in each individual account with the VPC in the infrastructure account,

D. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each subnet to associate with the resource share.

E. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each prefix list to associate with the resource share.

Answer

- B. **Enable resource sharing from the AWS Organizations management account.**
- D. **Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each subnet to associate with the resource share.**

Explanation

Enable AWS RAM sharing with AWS Organizations from the Organizations management account, then have the infrastructure account share the VPC subnets with the intended organizational unit. VPC subnet sharing lets the infrastructure team retain ownership and network-level management while accounts using the shared subnets can create their own AWS resources there. A transit gateway, duplicate VPCs, or shared prefix lists do not provide that subnet-sharing arrangement.

https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-sharing-share-subnet-working-with.html
https://docs.aws.amazon.com/en_en/vpc/latest/userguide/vpc-sharing.html
https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-ram.html

✗ **A** Create a transit gateway in the infrastructure account.
A transit gateway can connect VPCs and networks, but creating one does not share the infrastructure account's VPC subnets. It therefore does not by itself let other accounts create resources in those subnets while centralizing network management.

https://docs.aws.amazon.com/vpc/latest/userguide/vpc-sharing-share-subnet-working-with.html
https://docs.aws.amazon.com/en_en/vpc/latest/userguide/vpc-sharing.html

Free demo · full exam: 681 questions

✓ **B** Enable resource sharing from the AWS Organizations management account.
AWS RAM must be enabled for use with AWS Organizations so the infrastructure account can share resources with an organization or its organizational units without requiring individual invitations. This is the organization-level prerequisite for the subnet resource share in option D.

https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html
https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-ram.html
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-share-prerequisites.html

✗ **C** Create VPCs in each AWS account within the organization in AWS Organizations. Configure the VPCs to share the same CIDR range and subnets as the VPC in the infrastructure account. Peer the VPCs in each individual account with the VPC in the infrastructure account,
This creates separate VPCs rather than letting the individual accounts use subnets owned by the infrastructure account. VPC peering provides connectivity between VPCs, but it does not let accounts create resources in the infrastructure account's subnets or centralize ownership of their own VPCs.

https://docs.aws.amazon.com/vpc/latest/userguide/vpc-sharing-share-subnet-working-with.html
https://docs.aws.amazon.com/en_en/vpc/latest/userguide/vpc-sharing.html

✓ **D** Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each subnet to associate with the resource share.
The infrastructure account can use AWS RAM to share selected VPC subnets with an organizational unit. The infrastructure team retains control of the VPC and subnet configuration, while participating accounts can create supported resources in the shared subnets.

https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-sharing-share-subnet-working-with.html
https://docs.aws.amazon.com/en_en/vpc/latest/userguide/vpc-sharing.html

✗ **E** Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each prefix list to associate with the resource share.
Sharing prefix lists does not share the VPC subnets where the other accounts need to create resources. The resource share must include the subnets themselves to meet the requirement.

https://docs.aws.amazon.com/vpc/latest/userguide/vpc-sharing-share-subnet-working-with.html
https://docs.aws.amazon.com/ram/latest/userguide/shareable.html

Dumpsgate

Question 18

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design solutions that provide access to AWS accounts.

A company is using GitHub Actions to run a CI/CD pipeline that accesses resources on AWS. The company has an IAM user that uses a secret key in the pipeline to authenticate to AWS. An existing IAM role with an attached policy grants the required permissions to deploy resources.

The company's security team implements a new requirement that pipelines can no longer use longlived secret keys. A solutions architect must replace the secret key with a short-lived solution.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create an IAM SAML 2.0 identity provider (IdP) in IAM. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRole API call. Attach the existing IAM policy to the new IAM role. Update GitHub to use SAML authentication for the pipeline.

B. Create an IAM OpenID Connect (OIDC) identity provider (IdP) in IAM. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRoleWithWebIdentity API call from the GitHub OIDC IdP. Update GitHub to assume the role for the pipeline.

C. Create an Amazon Cognito identity pool. Configure the authentication provider to use GitHub. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRoleWithWebIdentity API call from the GitHub authentication provider. Configure the pipeline to use Cognito as its authentication provider.

D. Create a trust anchor to AWS Private CA. Generate a client certificate to use with AWS IAM Roles Anywhere. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRole API call. Attach the existing IAM policy to the new IAM role. Configure the pipeline to use the credential helper tool and to reference the client certificate public key to assume the new IAM role.

Answer: B. Create an IAM OpenID Connect (OIDC) identity provider (IdP) in IAM. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRoleWithWebIdentity API call from the GitHub OIDC IdP. Update GitHub to assume the role for the pipeline.

Explanation

GitHub Actions can present a short-lived OIDC token that AWS STS exchanges for temporary credentials by assuming an IAM role with `sts:AssumeRoleWithWebIdentity`. The new role can carry the existing deployment policy, so the pipeline no longer needs a long-lived IAM user secret key. Direct GitHub-to-IAM OIDC federation avoids the additional services and certificate management of the alternatives, meeting the least-operational-overhead requirement.

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_non-aws.html

✗ **A** Create an IAM SAML 2.0 identity provider (IdP) in IAM. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRole API call. Attach the existing IAM policy to the new IAM role. Update GitHub to use SAML authentication for the pipeline.
This uses SAML rather than the direct OIDC federation that GitHub Actions can use to obtain temporary AWS credentials. Also, the SAML role-assumption flow uses `sts:AssumeRoleWithSAML`, not the `sts:AssumeRole` call stated here. It is therefore not the appropriate low-overhead solution for this pipeline.

✓ **B** Create an IAM OpenID Connect (OIDC) identity provider (IdP) in IAM. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRoleWithWebIdentity API call from the GitHub OIDC IdP. Update GitHub to assume the role for the pipeline.
GitHub Actions can use OIDC to present a web identity token, which AWS STS exchanges for temporary

Free demo · full exam: 681 questions

role credentials through `sts:AssumeRoleWithWebIdentity`. The role can be given the existing deployment policy, eliminating the long-lived secret key without adding a separate credential service or certificate infrastructure.

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_non-aws.html

✗ **C** Create an Amazon Cognito identity pool. Configure the authentication provider to use GitHub. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRoleWithWebIdentity API call from the GitHub authentication provider. Configure the pipeline to use Cognito as its authentication provider.
An identity pool can provide temporary AWS credentials to authenticated identities, but adding Cognito as an intermediary is unnecessary when GitHub Actions can federate directly with IAM through OIDC. This adds configuration and operational components rather than minimizing overhead.

https://docs.aws.amazon.com/cognito/latest/developerguide/cognito-identity.html
https://docs.aws.amazon.com/cognitoidentity/latest/APIReference/API_CreateIdentityPool.html
https://docs.aws.amazon.com/cognito/latest/developerguide/getting-started-identity-pools-application.html

✗ **D** Create a trust anchor to AWS Private CA. Generate a client certificate to use with AWS IAM Roles Anywhere. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRole API call. Attach the existing IAM policy to the new IAM role. Configure the pipeline to use the credential helper tool and to reference the client certificate public key to assume the new IAM role.
IAM Roles Anywhere can issue temporary credentials to workloads outside AWS using X.509 certificates, but this approach requires a trust anchor, certificates, and credential-helper configuration. That certificate infrastructure creates more operational overhead than using GitHub Actions' direct OIDC integration.

https://docs.aws.amazon.com/rolesanywhere/latest/userguide/public-key-infrastructure.html
https://docs.aws.amazon.com/rolesanywhere/latest/userguide/introduction.html
https://docs.aws.amazon.com/rolesanywhere/latest/userguide/credential-helper.html
https://docs.aws.amazon.com/rolesanywhere/latest/userguide/trust-model.html

Dumpsgate

Question 19

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design solutions that provide access to AWS accounts.

Example Corp uses an organization in AWS Organizations with all features enabled to manage multiple AWS accounts. Example Corp acquires AnyCompany, which currently hosts its AWS environment in a single AWS account. Example Corp wants to add AnyCompany's AWS account to Example Corp's existing organization.

Which solution will meet this requirement?

A. From AnyCompany's AWS account, invite Example Corp's existing AWS management account to start the join process. Ensure that Organizations creates the OrganizationAccountAccessRole IAM role with the required permissions.

B. From Example Corp's management AWS account, send an AWS Support request to AWS to add AnyCompany's AWS account to the existing organization. Ensure that Organizations creates the OrganizationAccountAccessRole IAM role with the required permissions.

C. From Example Corp's management AWS account, send an invitation to AnyCompany's AWS account. Accept the invitation in AnyCompany's account. Create an OrganizationAccountAccessRole IAM role with the appropriate managed IAM policy.

D. From AnyCompany's AWS account, create an OrganizationAccountAccessRole IAM role with the appropriate managed IAM policy. Configure AnyCompany's AWS account to be a member of Example Corp's organization.

Answer: C. From Example Corp's management AWS account, send an invitation to AnyCompany's AWS account. Accept the invitation in AnyCompany's account. Create an OrganizationAccountAccessRole IAM role with the appropriate managed IAM policy.

Explanation

Because Example Corp already has an organization with all features enabled, its management account can invite AnyCompany's standalone AWS account to join. AnyCompany must accept the invitation, then create the OrganizationAccountAccessRole so Example Corp can access and manage the account across accounts. This is the sequence described in option C.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_create-cross-account-role.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_accept-decline-invite.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_invite-account.html

✗ **A** From AnyCompany's AWS account, invite Example Corp's existing AWS management account to start the join process. Ensure that Organizations creates the OrganizationAccountAccessRole IAM role with the required permissions.
The invitation must come from the organization, so Example Corp's management account should invite AnyCompany's account—not the other way around. AnyCompany is joining as a member account and cannot use this reversed invitation flow to add itself to Example Corp's organization.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_manage-invites.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_invite-account.html

✗ **B** From Example Corp's management AWS account, send an AWS Support request to AWS to add AnyCompany's AWS account to the existing organization. Ensure that Organizations creates the OrganizationAccountAccessRole IAM role with the required permissions.
Example Corp does not need AWS Support to add the account; the organization's management account can send an invitation directly. AnyCompany must accept that invitation to join, and the cross-account access role is handled separately.

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_accept-decline-invite.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_invite-account.html

✓ **C** From Example Corp's management AWS account, send an invitation to AnyCompany's AWS account. Accept the invitation in AnyCompany's account. Create an OrganizationAccountAccessRole IAM role with the appropriate managed IAM policy.
The organization's management account invites the standalone AnyCompany account, and AnyCompany accepts the invitation to become a member. Creating the OrganizationAccountAccessRole with the appropriate permissions then enables cross-account access from the organization's management account.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_create-cross-account-role.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_accept-decline-invite.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_invite-account.html

✗ **D** From AnyCompany's AWS account, create an OrganizationAccountAccessRole IAM role with the appropriate managed IAM policy. Configure AnyCompany's AWS account to be a member of Example Corp's organization.
Creating the role does not make AnyCompany's account a member of Example Corp's organization. The account must first be invited by the organization's management account and accept the invitation; configuring membership directly from AnyCompany is not the required join process.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_accept-decline-invite.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_invite-account.html

Question 20

Single choice · Domain: D1 Design Solutions for Organizational Complexity › Design architectures that are secure.

A company recently wanted a web application from an on-premises data center to the AWS Cloud.

The web application infrastructure consists of an Amazon CloudFront distribution that routes to an Application Load Balancer (ALB), with Amazon Elastic Container Service (Amazon ECS) to process requests. A recent security audit revealed that the web application is accessible by using both CloudFront and ALB endpoints. However. the company requires that the web application must be accessible only by using the CloudFront endpoint.

Which solution will meet this requirement with the LEAST amount of effort?

A. Create a new security group and attach it to the CloudFront distribution. Update the ALB security group ingress to allow access only from the CloudFront security group.

B. Update ALB security group ingress to allow access only from the CloudFront managed prefix list.

C. Create a VPC interface endpoint for Elastic Load Balancing. Update the ALB scheme from internet-facing to internal_

D. Extract CloudFront IPS from the AWS provided ip-ranges.json document. Update ALB security group ingress to allow access only from CloudFront IPs.

Answer: B. Update ALB security group ingress to allow access only from the CloudFront managed prefix list.

Explanation

B meets the requirement by restricting the ALB security group's inbound traffic to the AWS-managed prefix list for CloudFront origin-facing IP addresses. This allows CloudFront to reach the ALB while

blocking direct requests from other sources, including users connecting to the ALB endpoint. Using the managed prefix list avoids maintaining CloudFront IP addresses manually, making it the least-effort option.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html
https://docs.aws.amazon.com/vpc/latest/userguide/working-with-aws-managed-prefix-lists.html

✗ **A** Create a new security group and attach it to the CloudFront distribution. Update the ALB security group ingress to allow access only from the CloudFront security group.
A CloudFront distribution cannot have a security group attached to it, so the proposed security-group-to-security-group restriction cannot be configured as described. The relevant approach is to allow CloudFront origin-facing IP addresses in the ALB security group.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html

✓ **B** Update ALB security group ingress to allow access only from the CloudFront managed prefix list.
The CloudFront managed prefix list represents CloudFront origin-facing IP addresses and can be used as the source in an ALB security group ingress rule. This permits traffic from CloudFront while blocking direct access to the ALB from other sources, without requiring manual IP address maintenance.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html
https://docs.aws.amazon.com/vpc/latest/userguide/working-with-aws-managed-prefix-lists.html
https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/infrastructure-security.html

✗ **C** Create a VPC interface endpoint for Elastic Load Balancing. Update the ALB scheme from internet-facing to internal_
An interface endpoint for Elastic Load Balancing provides private access to the Elastic Load Balancing service API; it does not, by itself, configure CloudFront to reach an internal ALB. Changing the ALB to internal would require additional configuration and is not the least-effort solution for this requirement.

https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/load-balancer-vpc-endpoints.html

✗ **D** Extract CloudFront IPS from the AWS provided ip-ranges.json document. Update ALB security group ingress to allow access only from CloudFront IPs.
Allowing CloudFront IP ranges can restrict direct access, but extracting and maintaining those ranges from the published IP ranges document adds ongoing work as the ranges change. The CloudFront managed prefix list provides the same kind of source restriction with less maintenance.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html
https://docs.aws.amazon.com/vpc/latest/userguide/working-with-aws-managed-prefix-lists.html

Free demo · full exam: 681 questions

Question 21

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet reliability requirements.

A company has a collaborative document editing application that allows multiple users to simultaneously edit shared documents. The documents are stored as JSON objects in Amazon S3. The company needs to prevent data loss when multiple users modify the same document simultaneously. The solution must ensure that only the first write operation to finish successfully modifies the document. The solution must handle thousands of concurrent editors across multiple AWS Regions and must resolve conflicts automatically.

Which solution will meet these requirements?

A. Deploy an S3 bucket and enable versioning. Configure S3 event notifications to invoke an AWS Lambda function in response to PutObject events. Configure the function to track document versions in an Amazon DynamoDB table and to use conditional writes to resolve conflicts.

B. Deploy an S3 bucket. Configure conditional writes to use If-Match headers and entity tags (ETags). Configure an AWS Lambda function to detect editing conflicts based on Amazon S3 responses to writes. Configure the Lambda function to send retry messages to an Amazon SQS queue with a delay.

C. Deploy an S3 bucket and enable S3 Object Lock in compliance mode. Use AWS Step Functions to retry sequential document updates by using exponential backoff.

D. Deploy an S3 bucket and enable multipart uploads. Use AWS Batch jobs to make document updates. Use Amazon ElastiCache for Redis as a distributed lock mechanism across Regions.

Answer: B. Deploy an S3 bucket. Configure conditional writes to use If-Match headers and entity tags (ETags). Configure an AWS Lambda function to detect editing conflicts based on Amazon S3 responses to writes. Configure the Lambda function to send retry messages to an Amazon SQS queue with a delay.

Explanation

Option B uses an ETag with an If-Match conditional write so a write succeeds only if the document still has the version the editor read. When concurrent editors submit against the same ETag, the first successful write changes the ETag and later writes fail rather than silently overwriting it. The Lambda and SQS retry path can handle those conflicts by retrying the update, meeting the requirement to resolve conflicts automatically. [12]

https://docs.aws.amazon.com/us_en/AmazonS3/latest/userguide/conditional-writes.html

✗ **A** Deploy an S3 bucket and enable versioning. Configure S3 event notifications to invoke an AWS Lambda function in response to PutObject events. Configure the function to track document versions in an Amazon DynamoDB table and to use conditional writes to resolve conflicts.
S3 event notifications are triggered in response to object events, so they do not prevent a conflicting PutObject from modifying the document in the first place. Versioning can retain prior versions, but the proposed post-write notification and tracking do not provide the required conditional guard on the document write itself. [3]

https://docs.aws.amazon.com/AmazonS3/latest/userguide/EventNotifications.html

✓ **B** Deploy an S3 bucket. Configure conditional writes to use If-Match headers and entity tags (ETags). Configure an AWS Lambda function to detect editing conflicts based on Amazon S3 responses to writes. Configure the Lambda function to send retry messages to an Amazon SQS queue with a delay.
The If-Match condition makes the write conditional on the object still having the ETag the editor read. If another editor's write has already changed that ETag, the stale write fails instead of overwriting the

Free demo · full exam: 681 questions

newer document; the application can detect that conflict and retry through the queue. [12]

https://docs.aws.amazon.com/us_en/AmazonS3/latest/userguide/conditional-writes.html

✗ **C** Deploy an S3 bucket and enable S3 Object Lock in compliance mode. Use AWS Step Functions to retry sequential document updates by using exponential backoff.
Object Lock in compliance mode is intended to prevent an object version from being deleted or overwritten during its retention period, which conflicts with repeatedly editing shared documents. Sequential retries with backoff do not by themselves condition a write on the document version that the editor read. [5]

https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html

✗ **D** Deploy an S3 bucket and enable multipart uploads. Use AWS Batch jobs to make document updates. Use Amazon ElastiCache for Redis as a distributed lock mechanism across Regions.
Multipart upload is a way to upload an object in parts, not a mechanism for detecting that another editor has changed the document. Redis Global Datastore replicates data across Regions, but using it for a distributed lock does not provide the S3 write precondition needed to ensure stale document writes cannot overwrite a newer version. [10]

https://docs.aws.amazon.com/whitepapers/latest/database-caching-strategies-using-redis/amazon-elasticache-for-redis-global-datastore.html

Question 22

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet performance requirements.

A company runs an IoT application that consists of many connected IoT sensors. The company runs Eclipse Mosquitto MQTT brokers on two Amazon EC2 instances. The MQTT brokers send incoming data to a processing service that runs on a separate set of two EC2 instances that is the target for an Application Load Balancer (ALB). On average, the MQTT message rate is 1,500 messages each second. Surges of MQTT messages occur at unpredictable times throughout each day. Message surges can involve up to three times the average amount of traffic. The Mosquitto MQTT brokers cannot handle the increased number of connections.

The company needs a solution to improve the application's scalability.

Which solution will meet this requirement?

A. Migrate the MQTT brokers to AWS IoT Core. Create an AWS IoT rule to invoke an AWS Step Functions Express state machine to process data.

B. Migrate the MQTT brokers to AWS IoT Core. Create an AWS IoT rule to send incoming data to an Amazon EventBridge pipe. Invoke an AWS Lambda function from the pipe to process data.

C. Update the IoT sensors to send data to Amazon Data Firehose. Configure Firehose to invoke an AWS Lambda function to process data.

D. Update the IoT sensors to send data to an Amazon Kinesis data stream. Configure the data stream to invoke a new AWS Step Functions state machine to process data.

Answer: A. Migrate the MQTT brokers to AWS IoT Core. Create an AWS IoT rule to invoke an AWS Step Functions Express state machine to process data.

Explanation

AWS IoT Core replaces the self-managed Mosquitto brokers with a managed MQTT service, addressing

Free demo · full exam: 681 questions

the brokers' inability to handle connection surges. An AWS IoT rule can route incoming messages to a Step Functions state machine for processing, so the application can scale without relying on the two existing broker instances. This meets the requirement to improve scalability during unpredictable traffic increases.

https://docs.aws.amazon.com/iot/latest/developerguide/stepfunctions-rule-action.html
https://docs.aws.amazon.com/iot/latest/developerguide/iot-rule-actions.html
https://docs.aws.amazon.com/iot/latest/developerguide/iot-rules.html?trk=article-ssr-frontend-pulse_little-text-block

✓ **A** Migrate the MQTT brokers to AWS IoT Core. Create an AWS IoT rule to invoke an AWS Step Functions Express state machine to process data.
AWS IoT Core provides the managed MQTT broker layer, avoiding the connection-handling limit of the two Mosquitto instances. An IoT rule can route incoming messages to Step Functions, allowing processing to scale separately from the brokers.

https://docs.aws.amazon.com/iot/latest/developerguide/stepfunctions-rule-action.html
https://docs.aws.amazon.com/iot/latest/developerguide/iot-rule-actions.html
https://docs.aws.amazon.com/iot/latest/developerguide/iot-rules.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **B** Migrate the MQTT brokers to AWS IoT Core. Create an AWS IoT rule to send incoming data to an Amazon EventBridge pipe. Invoke an AWS Lambda function from the pipe to process data.
Although AWS IoT Core would address the broker connection limit, this option uses an EventBridge pipe as the IoT rule destination. The listed IoT rule actions do not provide a direct destination to an EventBridge pipe, so this proposed routing path does not meet the requirement as stated.

https://docs.aws.amazon.com/iot/latest/developerguide/iot-rule-actions.html

✗ **C** Update the IoT sensors to send data to Amazon Data Firehose. Configure Firehose to invoke an AWS Lambda function to process data.
This replaces the MQTT broker path with direct sensor delivery to Firehose, rather than providing a managed MQTT endpoint for the connected sensors. Firehose can invoke Lambda for data transformation, but that does not make it a replacement for the MQTT broker in this scenario.

https://docs.aws.amazon.com/firehose/latest/dev/create-transform.html

✗ **D** Update the IoT sensors to send data to an Amazon Kinesis data stream. Configure the data stream to invoke a new AWS Step Functions state machine to process data.
This requires changing the sensors to send to Kinesis instead of using the MQTT broker architecture. Also, a Kinesis data stream does not directly invoke a Step Functions state machine as described, so this option does not provide the stated processing path.

Free demo · full exam: 681 questions

Question 23

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet cost optimization goals.

A company uses an organization in AWS Organizations with all features enabled to manage AWS accounts. For each new project, the company creates a new linked account. After the creation of a new account, the root user signs in to the new account and creates a service request to increase the service quota for Amazon EC2 instances.

The company needs to automate this process.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Amazon EventBridge rule to detect the creation of a new account. Send the event to an Amazon SNS topic that invokes an AWS Lambda function. Configure the Lambda function to request a quota increase.

B. Create a Service Quotas request template in the organization's management account. Configure the request template to request a quota increase. Associate the request template with the organization.

C. Create an AWS Config rule in the organization's management account to check the quota. Create an AWS Lambda function. Configure the Lambda function to request a quota increase.

D. Create an automatic quota remediation action for AWS Trusted Advisor in the organization's management account. Create an AWS Step Functions workflow for the remediation action. Configure the Step Functions step to request a quota increase.

Answer: B. Create a Service Quotas request template in the organization's management account. Configure the request template to request a quota increase. Associate the request template with the organization.

Explanation

A Service Quotas request template lets the company define the EC2 quota increase in the Organizations management account and associate the template with the organization. The configured request is then submitted for accounts in the organization, including newly created linked accounts, without requiring the root user to sign in or a custom event-driven workflow. This meets the automation requirement with less operational overhead than building and maintaining custom components.

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html
https://docs.aws.amazon.com/pdfs/servicequotas/latest/userguide/ServiceQuotasUserGuide.pdf
https://docs.aws.amazon.com/servicequotas/latest/userguide/intro.html

✗ **A** Create an Amazon EventBridge rule to detect the creation of a new account. Send the event to an Amazon SNS topic that invokes an AWS Lambda function. Configure the Lambda function to request a quota increase.
This approach could automate a response to account creation, but it requires configuring and maintaining an event rule, SNS topic, and Lambda function. A Service Quotas request template provides the organization-level automation without that custom workflow.

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html
https://docs.aws.amazon.com/servicequotas/latest/userguide/eventbridge-integration.html
https://docs.aws.amazon.com/eventbridge/latest/ref/events-ref-organizations.html

✓ **B** Create a Service Quotas request template in the organization's management account. Configure the request template to request a quota increase. Associate the request template with the organization.
An organization-level Service Quotas request template allows the company to specify the quota increase request centrally and associate it with the organization. This automates requests for accounts in the organization, including new accounts, without requiring account-by-account root-user action or custom

Free demo · full exam: 681 questions

code.

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html
https://docs.aws.amazon.com/pdfs/servicequotas/latest/userguide/ServiceQuotasUserGuide.pdf
https://docs.aws.amazon.com/servicequotas/latest/userguide/intro.html

✗ **C** Create an AWS Config rule in the organization's management account to check the quota. Create an AWS Lambda function. Configure the Lambda function to request a quota increase.
AWS Config rules evaluate resource configurations; they are not the organization's built-in mechanism for submitting quota increase requests when accounts are created. Adding Lambda also creates custom automation to operate, whereas a Service Quotas request template directly addresses this scenario.

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html
https://docs.aws.amazon.com/us_en/config/latest/developerguide/evaluating-your-resources.html

✗ **D** Create an automatic quota remediation action for AWS Trusted Advisor in the organization's management account. Create an AWS Step Functions workflow for the remediation action. Configure the Step Functions step to request a quota increase.
This proposes a custom remediation workflow and therefore has more operational components to configure and maintain. The organization-level Service Quotas request template is the purpose-built option for automating quota requests across accounts.

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html
https://docs.aws.amazon.com/pdfs/servicequotas/latest/userguide/ServiceQuotasUserGuide.pdf
https://docs.aws.amazon.com/servicequotas/latest/userguide/intro.html

Question 24

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet cost optimization goals.

A company runs an application across three AWS accounts. The company uses an organization in AWS Organizations to manage the AWS accounts. One account hosts a development environment.

The second account hosts a testing environment. The third account hosts a production environment.

The application runs on Amazon ECS with the AWS Fargate launch type. The application stores data in an Amazon Aurora MySQL database.

The company expects production environment traffic to be consistent and predictable for the next year. The production environment remains active throughout each week. However, the development environment and the testing environment shut down during weekends and 12 hours each weekday.

The company wants to optimize costs for the application.

Which solution will meet this requirement?

A. In the organization management account, purchase Reserved Instances for a 1-year term for the databases. Purchase a 1-year EC2 Instance Savings Plan that will meet the compute usage demand.

B. In the organization management account, purchase Reserved Instances for a 1-year term for the production environment database. Purchase a 1-year Compute Savings Plan that will meet the compute usage demand.

C. In each account, purchase Reserved Instances for a 3-year term for the database. Purchase a 3-year Compute Savings Plan that will meet the compute usage demand.

D. In the production account, purchase Reserved Instances for a 3-year term for the database. Purchase a 3-year EC2 Instance Savings Plan that will meet the compute usage demand.

Free demo · full exam: 681 questions

Answer: B. In the organization management account, purchase Reserved Instances for a 1-year term for the production environment database. Purchase a 1-year Compute Savings Plan that will meet the compute usage demand.

Explanation

Option B matches the workloads that are expected to run consistently: a one-year Reserved Instance commitment for the production Aurora database and a one-year Compute Savings Plan for the application's compute usage. Compute Savings Plans apply to AWS Fargate, while EC2 Instance Savings Plans do not. The development and testing environments run only part of the time, so the company should avoid making long-term commitments sized for their intermittent usage; purchasing through the organization management account can allow eligible benefits to be shared across accounts.

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithReservedDBInstances.html
https://docs.aws.amazon.com/savingsplans/latest/userguide/plan-types.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html
https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-services.html

✗ **A** In the organization management account, purchase Reserved Instances for a 1-year term for the databases. Purchase a 1-year EC2 Instance Savings Plan that will meet the compute usage demand.
The EC2 Instance Savings Plan is not suitable for the application's Fargate compute; Compute Savings Plans cover Fargate. Also, buying database reservations for all three environments commits to development and testing database capacity even though those environments run only part of the time.

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithReservedDBInstances.html
https://docs.aws.amazon.com/savingsplans/latest/userguide/plan-types.html
https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-services.html
https://docs.aws.amazon.com/en_en/awsaccountbilling/latest/aboutv2/ri-turn-off.html

✓ **B** In the organization management account, purchase Reserved Instances for a 1-year term for the production environment database. Purchase a 1-year Compute Savings Plan that will meet the compute usage demand.
The production database has predictable, continuous usage, making a one-year Aurora Reserved Instance appropriate. A Compute Savings Plan can cover Fargate usage, and purchasing through the organization management account can make eligible discounts available across the organization's accounts.

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithReservedDBInstances.html
https://docs.aws.amazon.com/savingsplans/latest/userguide/plan-types.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html
https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-services.html

✗ **C** In each account, purchase Reserved Instances for a 3-year term for the database. Purchase a 3-year Compute Savings Plan that will meet the compute usage demand.
The development and testing environments are not continuously active, so database reservations sized for all accounts could go unused during their scheduled downtime while still creating a commitment. Three-year terms also exceed the stated one-year production traffic expectation.

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithReservedDBInstances.html
https://docs.aws.amazon.com/savingsplans/latest/userguide/plan-types.html
https://docs.aws.amazon.com/en_en/awsaccountbilling/latest/aboutv2/ri-turn-off.html

✗ **D** In the production account, purchase Reserved Instances for a 3-year term for the database. Purchase a 3-year EC2 Instance Savings Plan that will meet the compute usage demand.
A three-year commitment is longer than the stated one-year period of predictable production traffic. In addition, an EC2 Instance Savings Plan does not cover the application's Fargate compute; a Compute Savings Plan is the suitable plan type.

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithReservedDBInstances.html
https://docs.aws.amazon.com/savingsplans/latest/userguide/plan-types.html
https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-services.html

Free demo · full exam: 681 questions

Question 25

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet performance requirements.

A multinational insurance company is deploying an internal claims processing AI assistant that is powered by Amazon Bedrock FMs. The AI assistant must use several internal APIs to retrieve claim records, perform fraud checks, and update policy dat a. The company already exposes the APIs through a centralized Amazon Bedrock AgentCore gateway where enterprise policies enforce IAM authorization, request validation, and tool-level access controls. Some specialized validation logic for claims processing runs in containerized microservices.

During inference, the company wants the FMs to automatically discover and invoke only approved tools that are exposed through the enterprise integration layer. The company does not want to implement any application-side orchestration loop to run tools.

Which solution will meet these requirements?

A. Register AgentCore Runtime service endpoints as tools in the Amazon Bedrock Responses API so that Amazon Bedrock automatically discovers available tools that the company hosts in AgentCore Runtime.

B. Configure the Responses API to reference the AgentCore gateway ARN as the tool connector so the model retrieves and invokes the catalog of approved tools.

C. Establish an MCP connection between Amazon Bedrock and an AgentCore Runtime cluster to allow the AI assistant to invoke Runtime services.

D. Configure AgentCore Runtime services as Amazon Bedrock action groups. Invoke the action groups through the InvokeModelWithResponseStream API during FM inference.

Answer: B. Configure the Responses API to reference the AgentCore gateway ARN as the tool connector so the model retrieves and invokes the catalog of approved tools.

Explanation

The Responses API supports server-side tool execution with an Amazon Bedrock AgentCore Gateway, allowing the model to discover and invoke tools exposed through the gateway without the application implementing a tool-execution loop. Referencing the gateway as the tool connector uses the company's existing integration layer and its approved tool catalog, including the gateway's enterprise controls. This matches the requirement that the FM invoke only approved tools during inference. [1, 2, 3]

https://docs.aws.amazon.com/bedrock/latest/userguide/tool-use-server-side.html
https://aws.amazon.com/about-aws/whats-new/2026/02/amazon-bedrock-server-side-tool-execution-agentcore-gateway/
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/gateway-using.html

✗ **A** Register AgentCore Runtime service endpoints as tools in the Amazon Bedrock Responses API so that Amazon Bedrock automatically discovers available tools that the company hosts in AgentCore Runtime. This describes registering Runtime service endpoints, but the scenario's approved APIs are already exposed through an AgentCore Gateway. The Responses API's relevant server-side integration is with the gateway, which provides the tool catalog and enterprise integration controls; registering Runtime endpoints is not the specified way to use that catalog. [1, 2, 3]

https://docs.aws.amazon.com/bedrock/latest/userguide/tool-use-server-side.html
https://aws.amazon.com/about-aws/whats-new/2026/02/amazon-bedrock-server-side-tool-execution-agentcore-gateway/
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/gateway-using.html

✓ **B** Configure the Responses API to reference the AgentCore gateway ARN as the tool connector so the model retrieves and invokes the catalog of approved tools.

Free demo · full exam: 681 questions

The Responses API can use an AgentCore Gateway for server-side tool execution, so the model can discover and invoke tools exposed by that gateway without application-side orchestration. Using the existing gateway also keeps tool access within the enterprise integration layer and its policies. [1, 2, 3]

https://docs.aws.amazon.com/bedrock/latest/userguide/tool-use-server-side.html
https://aws.amazon.com/about-aws/whats-new/2026/02/amazon-bedrock-server-side-tool-execution-agentcore-gateway/
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/gateway-using.html

✗ **C** Establish an MCP connection between Amazon Bedrock and an AgentCore Runtime cluster to allow the AI assistant to invoke Runtime services.
The scenario calls for the FM to use the approved tools already exposed through the AgentCore Gateway, not to establish a separate connection to a Runtime cluster. A gateway can expose tools through MCP, but the Responses API solution here is to configure the gateway as the connector so its catalog and controls are used. [1, 2, 3, 5]

https://docs.aws.amazon.com/bedrock/latest/userguide/tool-use-server-side.html
https://aws.amazon.com/about-aws/whats-new/2026/02/amazon-bedrock-server-side-tool-execution-agentcore-gateway/
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/gateway-using.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/gateway-using-mcp-call.html

✗ **D** Configure AgentCore Runtime services as Amazon Bedrock action groups. Invoke the action groups through the InvokeModelWithResponseStream API during FM inference.
This uses Bedrock Agents action groups and an inference API rather than the Responses API's server-side gateway connector. It does not describe automatic discovery and server-side invocation of the gateway's approved tool catalog, which is the requirement in this scenario. [1, 8, 12]

https://docs.aws.amazon.com/bedrock/latest/userguide/tool-use-server-side.html
https://docs.aws.amazon.com/bedrock/latest/userguide/apis.html
https://docs.aws.amazon.com/bedrock/latest/userguide/inference-responses-api.html

Question 26

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet cost optimization goals.

A multi-tenant software as a service (SaaS) customer support platform serves thousands of enterprise clients. The platform must deploy more than 5,000 customized models that use the same ML framework to classify tickets, optimize inference infrastructure costs, and maintain sub-200 ms latency. The platform must operate in a multiaccount AWS architecture that includes separate accounts for development, staging, and production environments.

Which solution will meet these requirements MOST cost-effectively?

A. Deploy each model to a dedicated Amazon SageMaker AI real-time endpoint. Provision ml.g5.xlarge instances to handle peak traffic. Use cross-account IAM roles to grant AWS accounts access to model artifacts.

B. Deploy models by using Amazon SageMaker AI multi-model endpoints with inference components. Enable auto scaling in a centralized production account.

C. Deploy models to individual Amazon SageMaker AI serverless inference endpoints. Configure auto scaling. Deploy the endpoints in separate AWS accounts for each environment. Use cross-account Amazon S3 bucket policies to share models.

D. Deploy models to Amazon SageMaker AI asynchronous inference endpoints. Use Amazon S3 to queue requests from multiple accounts. Use batch transform jobs to perform ofline processing.

Free demo · full exam: 681 questions

Answer: B. Deploy models by using Amazon SageMaker AI multi-model endpoints with inference components. Enable auto scaling in a centralized production account.

Explanation

SageMaker AI multi-model endpoints let many models that use the same framework share endpoint infrastructure, avoiding the cost of provisioning a dedicated real-time endpoint for each of more than 5,000 customized models. Inference components and auto scaling let the production account scale serving capacity as demand changes while supporting real-time ticket classification. This is a better fit for the sub-200 ms latency target than asynchronous inference or batch processing.

https://docs.aws.amazon.com/sagemaker/latest/dg/multi-model-endpoints.html
https://docs.aws.amazon.com/sagemaker/latest/dg/multi-model-endpoints-autoscaling.html
https://docs.aws.amazon.com/sagemaker/latest/dg/batch-transform.html
https://docs.aws.amazon.com/sagemaker/latest/dg/async-inference.html

✗ **A** Deploy each model to a dedicated Amazon SageMaker AI real-time endpoint. Provision ml.g5.xlarge instances to handle peak traffic. Use cross-account IAM roles to grant AWS accounts access to model artifacts.
A dedicated endpoint for each of more than 5,000 models would duplicate serving infrastructure and incur substantial cost, especially when instances are sized for peak traffic. Cross-account access to model artifacts does not address that infrastructure cost. A shared multi-model endpoint is a more economical fit for models using the same framework.

https://docs.aws.amazon.com/sagemaker/latest/dg/multi-model-endpoints.html
https://docs.aws.amazon.com/sagemaker/latest/dg/realtime-endpoints-deploy-models.html

✓ **B** Deploy models by using Amazon SageMaker AI multi-model endpoints with inference components. Enable auto scaling in a centralized production account.
Multi-model endpoints allow many models to share endpoint infrastructure, reducing the cost of hosting thousands of customized models. Inference components and auto scaling support managing and scaling model serving in the production account, which fits the real-time classification and latency requirements better than asynchronous or batch approaches.

https://docs.aws.amazon.com/sagemaker/latest/dg/multi-model-endpoints.html
https://docs.aws.amazon.com/sagemaker/latest/dg/multi-model-endpoints-autoscaling.html

✗ **C** Deploy models to individual Amazon SageMaker AI serverless inference endpoints. Configure auto scaling. Deploy the endpoints in separate AWS accounts for each environment. Use cross-account Amazon S3 bucket policies to share models.
Creating an individual endpoint for each of thousands of models would retain substantial per-endpoint deployment overhead rather than sharing infrastructure. Serverless inference scales with demand, but it does not make this many separate endpoints the cost-effective choice, and its scale-from-zero behavior can introduce latency that conflicts with a sub-200 ms target. Cross-account S3 sharing does not resolve these issues.

https://docs.aws.amazon.com/sagemaker/latest/dg/serverless-endpoints.html
https://docs.aws.amazon.com/sagemaker/latest/dg/serverless-endpoints-autoscale.html

✗ **D** Deploy models to Amazon SageMaker AI asynchronous inference endpoints. Use Amazon S3 to queue requests from multiple accounts. Use batch transform jobs to perform ofline processing.
Asynchronous inference and batch transform are intended for workloads that can tolerate queued or offline processing, not interactive ticket classification with a sub-200 ms latency requirement. Using them would not meet the stated real-time serving need.

https://docs.aws.amazon.com/sagemaker/latest/dg/batch-transform.html
https://docs.aws.amazon.com/sagemaker/latest/dg/async-inference.html

Free demo · full exam: 681 questions

Question 27

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet reliability requirements.

A company deploys an AI agent on Amazon Bedrock AgentCore Runtime. The agent processes customer requests. Some agent actions are classified as high risk and must receive human approval before proceeding. The workflow must pause the agent, notify a human approver, and resume or cancel the action based on the approver's decision. Approvals typically take between 5 minutes and 60 minutes.

Which solution will meet these requirements?

A. Configure the agent to write high-risk action requests to an Amazon SQS queue. Create a separate polling application that reads the queue and sends approval email messages by using Amazon SES. Configure the application to invoke the agent by providing the approval result.

B. Configure an AWS Step Functions workflow that invokes the agent. Use a task token callback pattern to pause the workflow when the agent identifies a high-risk action. Send the task token to an approver by using Amazon SNS. Resume or cancel the workflow when the approver responds with the task token.

C. Configure the agent to invoke an AWS Lambda function for high-risk actions. Configure the Lambda function to send an approval email message by using Amazon SNS. Configure the Lambda function to poll an Amazon DynamoDB table until the approver updates the approval status.

D. Configure an Amazon EventBridge rule that intercepts all agent actions. Route high-risk actions to a separate approval queue. Create a second agent that monitors the queue and automatically approves or rejects actions without human input based on predefined rules.

Answer: B. Configure an AWS Step Functions workflow that invokes the agent. Use a task token callback pattern to pause the workflow when the agent identifies a high-risk action. Send the task token to an approver by using Amazon SNS. Resume or cancel the workflow when the approver responds with the task token.

Explanation

B is correct because a Step Functions task token callback pattern can pause the workflow while a high-risk action awaits a human decision. The workflow can notify the approver through SNS, then use the returned task token to continue after approval or stop after rejection. This directly supports the required human approval and pause/resume behavior for decisions that typically take 5 to 60 minutes.

https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-human-approval.html
https://docs.aws.amazon.com/step-functions/latest/dg/connect-to-resource.html
https://docs.aws.amazon.com/step-functions/latest/dg/connect-sns.html

✗ **A** Configure the agent to write high-risk action requests to an Amazon SQS queue. Create a separate polling application that reads the queue and sends approval email messages by using Amazon SES. Configure the application to invoke the agent by providing the approval result.
This design uses a separate polling application rather than a workflow callback to suspend the specific action and resume or cancel it with an approval response. It also adds custom coordination to connect the approver's decision back to the agent.

✓ **B** Configure an AWS Step Functions workflow that invokes the agent. Use a task token callback pattern to pause the workflow when the agent identifies a high-risk action. Send the task token to an approver by using Amazon SNS. Resume or cancel the workflow when the approver responds with the task token.
The task token callback pattern lets Step Functions pause at the high-risk action until an external party responds, then continue or fail the workflow based on that response. SNS can be used to notify the approver, and the decision can be associated with the paused task through its token.

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-human-approval.html
https://docs.aws.amazon.com/step-functions/latest/dg/connect-to-resource.html
https://docs.aws.amazon.com/step-functions/latest/dg/connect-sns.html

✗ **C** Configure the agent to invoke an AWS Lambda function for high-risk actions. Configure the Lambda function to send an approval email message by using Amazon SNS. Configure the Lambda function to poll an Amazon DynamoDB table until the approver updates the approval status.
Polling DynamoDB inside a Lambda function is not the callback-based pause-and-resume design described by the requirements. It also keeps the function responsible for waiting and checking for a decision rather than letting an orchestrated workflow pause while the human approval is pending.

✗ **D** Configure an Amazon EventBridge rule that intercepts all agent actions. Route high-risk actions to a separate approval queue. Create a second agent that monitors the queue and automatically approves or rejects actions without human input based on predefined rules.
The scenario requires a human approver to make the decision. A second agent that automatically approves or rejects actions does not provide human approval, even if it can route high-risk actions to a separate queue.

Question 28

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet performance requirements.

A retail company has an ecommerce application that uses ML to make product recommendations.

The company creates two new ML model variants to improve the recommendations. The company wants to A/B test the two model variants to determine which one performs better. The company wants to split traffic evenly between the two variants. The company wants a solution that minimizes the impact to the existing ecommerce application.

Which solution will meet these requirements?

A. Deploy both model variants to an Amazon SageMaker AI serverless inference endpoint. In the endpoint configuration, set the weight for each model variant to 0.5. In the ecommerce application, invoke the inference endpoint.

B. Deploy each model variant to its own Amazon SageMaker AI serverless inference endpoint. In the ecommerce application, invoke each model variant endpoint directly by using a round robin strategy. Log variant invocations in Amazon CloudWatch Logs.

C. Deploy both variants to an Amazon SageMaker AI real-time inference endpoint. In the endpoint configuration, set the weight for each model variant to 0.5. In the ecommerce application, invoke the inference endpoint.

D. Deploy both variants to an Amazon SageMaker AI real-time inference endpoint. In the ecommerce application, invoke each model variant directly by using a round robin strategy. Log variant invocations in Amazon CloudWatch Logs.

Answer: C. Deploy both variants to an Amazon SageMaker AI real-time inference endpoint. In the endpoint configuration, set the weight for each model variant to 0.5. In the ecommerce application, invoke the inference endpoint.

Explanation

A SageMaker AI real-time inference endpoint can host both model variants as production variants and route requests between them using configured variant weights. Setting each weight to 0.5 splits traffic

evenly, while the ecommerce application continues invoking a single endpoint rather than implementing its own routing. This minimizes changes to the existing application and supports the requested A/B test.

https://docs.aws.amazon.com/sagemaker/latest/dg/model-ab-testing.html
https://docs.aws.amazon.com/sagemaker/latest/APIReference/API_ProductionVariant.html

✗ **A** Deploy both model variants to an Amazon SageMaker AI serverless inference endpoint. In the endpoint configuration, set the weight for each model variant to 0.5. In the ecommerce application, invoke the inference endpoint.
The approach of assigning weights to production variants is suitable for a real-time endpoint, but serverless inference endpoints do not support this multi-variant weighted traffic-splitting setup. Therefore, this option cannot provide the required A/B test as described.

https://docs.aws.amazon.com/sagemaker/latest/dg/model-ab-testing.html
https://docs.aws.amazon.com/sagemaker/latest/APIReference/API_ProductionVariant.html
https://docs.aws.amazon.com/sagemaker/latest/dg/serverless-endpoints-autoscale.html

✗ **B** Deploy each model variant to its own Amazon SageMaker AI serverless inference endpoint. In the ecommerce application, invoke each model variant endpoint directly by using a round robin strategy. Log variant invocations in Amazon CloudWatch Logs.
The application would have to call two endpoints and implement the round-robin routing itself, rather than sending requests to one endpoint that handles variant allocation. That adds application changes and does not minimize impact to the existing ecommerce application.

https://docs.aws.amazon.com/sagemaker/latest/dg/model-ab-testing.html
https://docs.aws.amazon.com/sagemaker/latest/APIReference/API_runtime_InvokeEndpoint.html

✓ **C** Deploy both variants to an Amazon SageMaker AI real-time inference endpoint. In the endpoint configuration, set the weight for each model variant to 0.5. In the ecommerce application, invoke the inference endpoint.
A real-time endpoint can host multiple production variants and use configured variant weights to distribute requests. Equal weights split traffic evenly, and the application only needs to invoke the endpoint, minimizing changes to the existing application.

https://docs.aws.amazon.com/sagemaker/latest/dg/model-ab-testing.html
https://docs.aws.amazon.com/sagemaker/latest/APIReference/API_ProductionVariant.html

✗ **D** Deploy both variants to an Amazon SageMaker AI real-time inference endpoint. In the ecommerce application, invoke each model variant directly by using a round robin strategy. Log variant invocations in Amazon CloudWatch Logs.
Although the variants can be deployed to a real-time endpoint, having the application select each variant in a round-robin strategy puts traffic allocation in the application. This requires more application changes than configuring weights on the endpoint, so it does not best meet the requirement to minimize impact.

https://docs.aws.amazon.com/sagemaker/latest/dg/model-ab-testing.html
https://docs.aws.amazon.com/sagemaker/latest/APIReference/API_runtime_InvokeEndpoint.html

Free demo · full exam: 681 questions

Question 29

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet performance requirements.

A machine learning (ML) company continuously develops new versions of its models. The models are integrated into a web application that uses a RESTful API to call custom Python code to invoke the models. The application runs in an Amazon EC2 Auto Scaling group.

Customer feedback indicates that new model versions sometimes perform worse than previous versions. The company wants to implement a gradual deployment approach in which 10% of traffic is initially routed to new models to verify performance. The solution must also reduce operational overhead for infrastructure management.

Which solution will meet these requirements with the LEAST operational overhead?

A. Deploy model versions to an Amazon SageMaker AI multi-model endpoint. Configure the weight for the new version to 10% and the previous version to 90%. Migrate the Python code to an AWS Lambda function. Use Amazon API Gateway to invoke the Lambda function.

B. Deploy each model version to a separate Amazon SageMaker Serverless Inference endpoint. Configure an AWS Lambda function to route 10% of invocation calls to the new endpoint and 90% to the previous version based on weighted random selection. Use Amazon API Gateway to invoke the Lambda function.

C. Deploy each model version to a separate Amazon SageMaker AI single-model endpoint. Configure an AWS Lambda function to route 10% of invocation calls to the new endpoint's URL and 90% to the previous version. Use Amazon API Gateway to invoke the Lambda function.

D. Deploy model versions to an Amazon SageMaker Serverless Inference endpoint. Use Amazon API Gateway with a canary deployment to gradually shift 10% of traffic to the new model version while maintaining 90% of traffic on the previous version.

Answer: B. Deploy each model version to a separate Amazon SageMaker Serverless Inference endpoint. Configure an AWS Lambda function to route 10% of invocation calls to the new endpoint and 90% to the previous version based on weighted random selection. Use Amazon API Gateway to invoke the Lambda function.

Explanation

B uses separate SageMaker Serverless Inference endpoints for the old and new models, avoiding the need to manage endpoint instances. A Lambda function can send 10% of calls to the new endpoint and 90% to the previous endpoint, enabling a gradual performance check. API Gateway invokes the Lambda function, replacing the existing EC2 Auto Scaling application path with a lower-operations serverless design. [2]

https://docs.aws.amazon.com/sagemaker/latest/dg/serverless-endpoints.html

✗ **A** Deploy model versions to an Amazon SageMaker AI multi-model endpoint. Configure the weight for the new version to 10% and the previous version to 90%. Migrate the Python code to an AWS Lambda function. Use Amazon API Gateway to invoke the Lambda function.
A multi-model endpoint can host multiple models, but the stated model-level 10%/90% weight configuration does not by itself route traffic between models hosted on that endpoint. Production-variant weights apply to variants, so this option does not provide the specified routing as written. [4] [12]

https://docs.aws.amazon.com/sagemaker/latest/APIReference/API_ProductionVariant.html
https://docs.aws.amazon.com/sagemaker/latest/dg/build-multi-model-build-container.html

✓ **B** Deploy each model version to a separate Amazon SageMaker Serverless Inference endpoint. Configure an AWS Lambda function to route 10% of invocation calls to the new endpoint and 90% to the

Free demo · full exam: 681 questions

previous version based on weighted random selection. Use Amazon API Gateway to invoke the Lambda function.
The Lambda function explicitly applies the required 10%/90% split between the two model endpoints. SageMaker Serverless Inference avoids managing endpoint instances, while API Gateway and Lambda provide the request path without the existing EC2 Auto Scaling application infrastructure. [2]

https://docs.aws.amazon.com/sagemaker/latest/dg/serverless-endpoints.html

✗ **C** Deploy each model version to a separate Amazon SageMaker AI single-model endpoint. Configure an AWS Lambda function to route 10% of invocation calls to the new endpoint's URL and 90% to the previous version. Use Amazon API Gateway to invoke the Lambda function.
The Lambda routing would implement the desired traffic split, but separate single-model endpoints require managing their endpoint infrastructure. Serverless Inference provides a lower-operations way to host the versions, so this does not best meet the least-operational-overhead requirement. [2]

https://docs.aws.amazon.com/sagemaker/latest/dg/serverless-endpoints.html

✗ **D** Deploy model versions to an Amazon SageMaker Serverless Inference endpoint. Use Amazon API Gateway with a canary deployment to gradually shift 10% of traffic to the new model version while maintaining 90% of traffic on the previous version.
API Gateway canary deployments split traffic between API deployments; they do not themselves select between two model versions hosted behind a single SageMaker endpoint. The option does not specify separate model destinations or another mechanism for that model-level routing. [1] [3]

https://docs.aws.amazon.com/apigateway/latest/api/API_DeploymentCanarySettings.html
https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html

Dumpsgate

Question 30

Single choice · Domain: D2 Design for New Solutions › Determine a solution design to meet business requirements.

A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.

The company wants to stop using email in the application and move to a serverless application model.

Which architecture solution meets these requirements?

A. Use AWS Batch to configure the different tasks required to ship a package. Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping label. Once that label is scanned, as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.

B. When a new order is created, store the order information in Amazon SQS. Have AWS Lambda check the queue every 5 minutes and process any needed work. When an order needs to be shipped, have Lambda print the label in the warehouse. Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon SQS.

C. Update the application to store new order information in Amazon DynamoDB. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress," and print a package label to the warehouse. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.

D. Store new order information in Amazon EFS. Have instances pull the new information from the NFS and send that information to printers in the warehouse. Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.

Answer: C. Update the application to store new order information in Amazon DynamoDB. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress," and print a package label to the warehouse. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.

Explanation

Option C replaces the email-based process with an event-driven workflow: order data is stored in DynamoDB, and a new order starts a Step Functions workflow to coordinate the in-progress and shipping-label steps. After the label is scanned and fulfillment is confirmed, Lambda can update the order to shipped and complete the workflow. Step Functions and Lambda provide a serverless way to coordinate these application steps. [1][2]

https://docs.aws.amazon.com/lambda/latest/dg/with-step-functions.html
https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-creating-lambda-state-machine.html

✗ **A** Use AWS Batch to configure the different tasks required to ship a package. Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping label. Once that label is scanned, as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.
AWS Batch is intended to run batch computing jobs, whereas this scenario needs to coordinate an order's ongoing steps and react when its label is scanned. The option does not provide an appropriate serverless workflow for tracking that order from processing through shipment.

Free demo · full exam: 681 questions

✗ **B** When a new order is created, store the order information in Amazon SQS. Have AWS Lambda check the queue every 5 minutes and process any needed work. When an order needs to be shipped, have Lambda print the label in the warehouse. Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon SQS.
The design relies on an EC2 instance to update the queue, so it does not move the application to a serverless model. It also has Lambda check the queue on a five-minute schedule rather than using a queue-triggered Lambda integration; AWS documents event source mappings for processing SQS messages with Lambda. [4][5][11]

https://docs.aws.amazon.com/lambda/latest/dg/services-sqs-configure.html
https://docs.aws.amazon.com/lambda/latest/dg/with-sqs.html
https://docs.aws.amazon.com/lambda/latest/dg/invocation-eventsourcemapping.html

✓ **C** Update the application to store new order information in Amazon DynamoDB. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress," and print a package label to the warehouse. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.
This option uses a serverless workflow to coordinate the order's steps, from marking it in progress and printing a label to recording shipment after fulfillment. Step Functions can orchestrate Lambda functions, making this a suitable replacement for the email-based process. [1][2]

https://docs.aws.amazon.com/lambda/latest/dg/with-step-functions.html
https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-creating-lambda-state-machine.html

✗ **D** Store new order information in Amazon EFS. Have instances pull the new information from the NFS and send that information to printers in the warehouse. Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.
This design depends on instances to poll EFS and handle the warehouse tasks, so it retains server-managed compute rather than adopting a serverless application model. It also lacks a workflow service to coordinate the order's processing and completion steps.

Question 31

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet reliability requirements.

A company is building an automated pipeline that uses AWS Step Functions to orchestrate multiple stages. The pipeline includes a confirmation stage.

The company wants to avoid unnecessary compute consumption. The company must ensure that the workflow can remain pending over extended periods without state loss. The confirmation signal originates from a control-plane service that runs in a separate AWS account and is emitted asynchronously after human review. The solution must ensure that only the authorized external system can resume execution. The solution should not expose long-lived credentials.

Which solution will meet these requirements?

A. Configure a Wait state with dynamically extended durations. Invoke an AWS Lambda function that checks approval status from the external system on each retry attempt.

B. Configure a Task state that uses a callback pattern to provide a unique correlation token and to resume execution only when the external system returns that token after approval.

C. Configure an activity worker that retrieves tasks and holds the workflow execution open within its own long-running process until approval is observed externally.

D. Configure an Amazon EventBridge rule that periodically triggers execution resumption by reinvoking the workflow with inputs that reflect approval state and correlation identifiers.

Free demo · full exam: 681 questions

Answer: B. Configure a Task state that uses a callback pattern to provide a unique correlation token and to resume execution only when the external system returns that token after approval.

Explanation

A callback Task state can pause the workflow while human review happens asynchronously, without repeatedly running compute to check for approval. Step Functions supplies a unique task token, and the authorized external system can return that token through the callback API to resume the waiting task. The external system can use appropriately scoped cross-account IAM access rather than long-lived credentials, and the workflow retains its pending state while it waits.

https://docs.aws.amazon.com/step-functions/latest/dg/concepts-access-cross-acct-resources.html
https://docs.aws.amazon.com/step-functions/latest/dg/connect-to-resource.html
https://docs.aws.amazon.com/step-functions/latest/apireference/API_SendTaskSuccess.html
https://docs.aws.amazon.com/step-functions/latest/dg/security_iam_service-with-iam.html

✗ **A** Configure a Wait state with dynamically extended durations. Invoke an AWS Lambda function that checks approval status from the external system on each retry attempt.
This design repeatedly invokes Lambda to poll for approval, which consumes compute even when no decision has been made. A Wait state can pause execution for a specified duration, but polling is not the asynchronous, token-based confirmation mechanism required here.

https://docs.aws.amazon.com/step-functions/latest/dg/state-wait.html

✓ **B** Configure a Task state that uses a callback pattern to provide a unique correlation token and to resume execution only when the external system returns that token after approval.
The callback pattern lets the Task wait for an asynchronous approval without running a polling process, and the task token correlates the approval with the specific waiting execution. The authorized external service can call the callback API with that token to resume the task, using scoped IAM permissions for cross-account access rather than long-lived credentials.

https://docs.aws.amazon.com/step-functions/latest/dg/concepts-access-cross-acct-resources.html
https://docs.aws.amazon.com/step-functions/latest/dg/connect-to-resource.html
https://docs.aws.amazon.com/step-functions/latest/apireference/API_SendTaskSuccess.html
https://docs.aws.amazon.com/step-functions/latest/dg/security_iam_service-with-iam.html

✗ **C** Configure an activity worker that retrieves tasks and holds the workflow execution open within its own long-running process until approval is observed externally.
An activity worker must run and poll for work, and the described long-running process continues consuming compute while approval is pending. That conflicts with the requirement to avoid unnecessary compute; a callback Task can leave the workflow waiting without such a worker process.

https://docs.aws.amazon.com/step-functions/latest/dg/concepts-activities.html

✗ **D** Configure an Amazon EventBridge rule that periodically triggers execution resumption by reinvoking the workflow with inputs that reflect approval state and correlation identifiers.
Periodically triggering a workflow is a polling-style approach and does not use the callback token to resume the particular waiting Task. Reinvoking the workflow with approval inputs does not provide the direct, token-correlated callback required by the scenario.

https://docs.aws.amazon.com/step-functions/latest/dg/connect-to-resource.html
https://docs.aws.amazon.com/step-functions/latest/apireference/API_SendTaskSuccess.html

Free demo · full exam: 681 questions

Question 32

Multiple choice · Domain: D2 Design for New Solutions › Design a solution to meet reliability requirements.

A company needs to store sensitive and crucial business data in an Amazon S3 bucket. Because of the criticality of the data, the data must be available in a second AWS Region within 15 minutes. If any object is not available within 15 minutes, an operations team must receive a notification. The company has already enabled S3 Versioning on the S3 bucket, and the company creates a second S3 bucket in the second Region.

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

A. Create a new S3 replication rule. Configure the rule to replicate all objects to the second S3 bucket.

B. Create an AWS DataSync task. Set the appropriate source and destination S3 buckets for the task.

C. Enable S3 Replication Time Control (S3 RTC), and configure the replication metrics and events. Create an Amazon SNS topic and set it as the target for the OperationMissedThreshold S3 event notification. Use the SNS topic to notify the operations team.

D. Enable S3 Transfer Acceleration on both S3 buckets. Create an Amazon CloudWatch alarm for the S3 5xxErrors metric. Create an Amazon SNS topic and set it as the target for the alarm. Use the SNS topic to notify the operations team.

E. Enable AWS DataSync enhanced transfer mode. Create an Amazon CloudWatch alarm for the DataSync FilesFailed metric. Create an Amazon SNS topic and set it as the target for the alarm. Use the SNS topic to notify the operations team.

Answer

- A. **Create a new S3 replication rule. Configure the rule to replicate all objects to the second S3 bucket.**
- C. **Enable S3 Replication Time Control (S3 RTC), and configure the replication metrics and events. Create an Amazon SNS topic and set it as the target for the OperationMissedThreshold S3 event notification. Use the SNS topic to notify the operations team.**

Explanation

A cross-Region S3 replication rule copies objects from the source bucket to the second-Region bucket. S3 Replication Time Control (S3 RTC) is designed to replicate most new objects within 15 minutes and provides replication metrics and events to track objects that miss that threshold. Configuring the OperationMissedThreshold event to notify an SNS topic alerts the operations team when replication exceeds the required time.

https://docs.aws.amazon.com/en_en/AmazonS3/latest/userguide/replication-time-control.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-metrics-events.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-how-setup.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-add-config.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html

✓ **A** Create a new S3 replication rule. Configure the rule to replicate all objects to the second S3 bucket. The replication rule sends objects from the source bucket to the destination bucket in the second Region, meeting the cross-Region availability requirement. Configure the rule to cover all objects; S3 replication requires versioning on both buckets.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-how-setup.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-add-config.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html

Free demo · full exam: 681 questions

✗ **B** Create an AWS DataSync task. Set the appropriate source and destination S3 buckets for the task.
A DataSync task can transfer data between S3 buckets, but the option does not provide the required 15-minute replication commitment or notification when an object misses that threshold. S3 RTC with replication events addresses both requirements directly.

https://docs.aws.amazon.com/en_en/AmazonS3/latest/userguide/replication-time-control.html
https://docs.aws.amazon.com/datasync/latest/userguide/task-scheduling.html

✓ **C** Enable S3 Replication Time Control (S3 RTC), and configure the replication metrics and events. Create an Amazon SNS topic and set it as the target for the OperationMissedThreshold S3 event notification. Use the SNS topic to notify the operations team.
S3 RTC is designed to replicate 99.99% of new objects within 15 minutes. Its replication metrics and events can identify operations that miss the threshold, and routing the OperationMissedThreshold notification to SNS alerts the operations team.

https://docs.aws.amazon.com/en_en/AmazonS3/latest/userguide/replication-time-control.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-metrics-events.html

✗ **D** Enable S3 Transfer Acceleration on both S3 buckets. Create an Amazon CloudWatch alarm for the S3 5xxErrors metric. Create an Amazon SNS topic and set it as the target for the alarm. Use the SNS topic to notify the operations team.
Transfer Acceleration is not a cross-Region replication mechanism and does not ensure objects reach the second bucket within 15 minutes. A 5xxErrors alarm detects server errors, not whether replication has missed its time threshold.

✗ **E** Enable AWS DataSync enhanced transfer mode. Create an Amazon CloudWatch alarm for the DataSync FilesFailed metric. Create an Amazon SNS topic and set it as the target for the alarm. Use the SNS topic to notify the operations team.
A FilesFailed alarm would report failed file transfers, not objects that have not replicated within 15 minutes. This setup also does not provide the S3 RTC replication-time commitment required by the scenario.

https://docs.aws.amazon.com/en_en/AmazonS3/latest/userguide/replication-time-control.html
https://docs.aws.amazon.com/datasync/latest/userguide/task-scheduling.html

Dumpsgate

Free demo · full exam: 681 questions

Question 33

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet cost optimization goals.

A company that has multiple AWS accounts is using AWS Organizations. The company's AWS accounts host VPCs, Amazon EC2 instances, and containers.

The company's compliance team has deployed a security tool in each VPC where the company has deployments. The security tools run on EC2 instances and send information to the AWS account that is dedicated for the compliance team. The company has tagged all the compliance-related resources with a key of "costCenter" and a value or "compliance".

The company wants to identify the cost of the security tools that are running on the EC2 instances so that the company can charge the compliance team's AWS account. The cost calculation must be as accurate as possible.

What should a solutions architect do to meet these requirements?

A. In the management account of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.

B. In the member accounts of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Schedule a monthly AWS Lambda function to retrieve the reports and calculate the total cost for the costCenter tagged resources.

C. In the member accounts of the organization activate the costCenter user-defined tag. From the management account, schedule a monthly AWS Cost and Usage Report. Use the tag breakdown in the report to calculate the total cost for the costCenter tagged resources.

D. Create a custom report in the organization view in AWS Trusted Advisor. Configure the report to generate a monthly billing summary for the costCenter tagged resources in the compliance team's AWS account.

Answer: A. In the management account of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.

Explanation

Activate the user-defined costCenter tag in the organization's management account so it can be used for cost allocation across the organization's accounts. A monthly Cost and Usage Report delivered to an S3 bucket in that account can include the tag breakdown, allowing the company to identify the costs of the tagged security-tool resources and charge them to the compliance team. This uses consolidated billing data and avoids a separate custom calculation.

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/account-tags-cost-allocation.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html
https://docs.aws.amazon.com/en_en/awsaccountbilling/latest/aboutv2/configurecostallocreport.html
https://docs.aws.amazon.com/cur/latest/userguide/cur-consolidated-billing.html

✓ **A** In the management account of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.
The management account is where cost allocation tags are activated for use in the organization's

Free demo · full exam: 681 questions

consolidated billing data. The Cost and Usage Report can include the activated tag and provide a breakdown of costs for resources tagged costCenter=compliance across the accounts, supporting the requested chargeback.

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/account-tags-cost-allocation.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html
https://docs.aws.amazon.com/en_en/awsaccountbilling/latest/aboutv2/configurecostallocreport.html
https://docs.aws.amazon.com/cur/latest/userguide/cur-consolidated-billing.html

✗ **B** In the member accounts of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Schedule a monthly AWS Lambda function to retrieve the reports and calculate the total cost for the costCenter tagged resources.
Activating the tag in the member accounts is not the organization-level cost allocation setup needed for consolidated billing; the tag should be activated in the management account. The Lambda calculation is also unnecessary because the Cost and Usage Report can provide the cost breakdown for the activated tag.

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/account-tags-cost-allocation.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html
https://docs.aws.amazon.com/en_en/awsaccountbilling/latest/aboutv2/configurecostallocreport.html
https://docs.aws.amazon.com/cur/latest/userguide/cur-consolidated-billing.html

✗ **C** In the member accounts of the organization activate the costCenter user-defined tag. From the management account, schedule a monthly AWS Cost and Usage Report. Use the tag breakdown in the report to calculate the total cost for the costCenter tagged resources.
Although the management account is the appropriate place to configure the organization's Cost and Usage Report, activating the tag only in member accounts does not meet the cost-allocation-tag requirement. The tag must be activated in the management account for the report's tag breakdown to support this consolidated cost analysis.

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/account-tags-cost-allocation.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html
https://docs.aws.amazon.com/en_en/awsaccountbilling/latest/aboutv2/configurecostallocreport.html
https://docs.aws.amazon.com/cur/latest/userguide/cur-consolidated-billing.html

✗ **D** Create a custom report in the organization view in AWS Trusted Advisor. Configure the report to generate a monthly billing summary for the costCenter tagged resources in the compliance team's AWS account.
Trusted Advisor organization reports are not the mechanism for producing a monthly billing summary broken down by user-defined cost allocation tags. A Cost and Usage Report with the tag activated in the management account is the appropriate source for calculating these costs.

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/account-tags-cost-allocation.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html
https://docs.aws.amazon.com/en_en/awsaccountbilling/latest/aboutv2/configurecostallocreport.html

Free demo · full exam: 681 questions

Question 34

Single choice · Domain: D2 Design for New Solutions › Determine a solution design to meet business requirements.

A software company needs to create short-lived test environments to test pull requests as part of its development process. Each test environment consists of a single Amazon EC2 instance that is in an Auto Scaling group.

The test environments must be able to communicate with a central server to report test results. The central server is located in an on-premises data center. A solutions architect must implement a solution so that the company can create and delete test environments without any manual intervention. The company has created a transit gateway with a VPN attachment to the on-premises network.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS CloudFormation template that contains a transit gateway attachment and related routing configurations. Create a CloudFormation stack set that includes this template. Use CloudFormation StackSets to deploy a new stack for each VPC in the account. Deploy a new VPC for each test environment.

B. Create a single VPC for the test environments. Include a transit gateway attachment and related routing configurations. Use AWS CloudFormation to deploy all test environments into the VPC.

C. Create a new OU in AWS Organizations for testing. Create an AWS CloudFormation template that contains a VPC, necessary networking resources, a transit gateway attachment, and related routing configurations. Create a CloudFormation stack set that includes this template. Use CloudFormation StackSets for deployments into each account under the testing 01.1. Create a new account for each test environment.

D. Convert the test environment EC2 instances into Docker images. Use AWS CloudFormation to configure an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in a new VPC, create a transit gateway attachment, and create related routing configurations. Use Kubernetes to manage the deployment and lifecycle of the test environments.

Answer: B. Create a single VPC for the test environments. Include a transit gateway attachment and related routing configurations. Use AWS CloudFormation to deploy all test environments into the VPC.

Explanation

A single VPC can host the short-lived EC2 test environments, with one transit gateway attachment and the required routes configured for that VPC. CloudFormation can create and delete the environments in that VPC without repeating the attachment and routing setup for each environment. This lets the instances reach the on-premises central server through the existing transit gateway and VPN while minimizing operational overhead.

https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpc-attachments.html
https://docs.aws.amazon.com/vpc/latest/tgw/how-transit-gateways-work.html
https://docs.aws.amazon.com/en_en/vpc/latest/userguide/route-table-options.html

✗ **A** Create an AWS CloudFormation template that contains a transit gateway attachment and related routing configurations. Create a CloudFormation stack set that includes this template. Use CloudFormation StackSets to deploy a new stack for each VPC in the account. Deploy a new VPC for each test environment.
Creating a VPC and transit gateway attachment for every short-lived test environment repeats networking setup that can be shared. This adds operational work and resources to manage and remove,

Free demo · full exam: 681 questions

rather than using one VPC and attachment for all test environments.

https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpc-attachments.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html
https://docs.aws.amazon.com/vpc/latest/tgw/how-transit-gateways-work.html

✓ **B** Create a single VPC for the test environments. Include a transit gateway attachment and related routing configurations. Use AWS CloudFormation to deploy all test environments into the VPC.
The shared VPC needs only one transit gateway attachment and its routing configured for connectivity to the on-premises network. CloudFormation can then deploy the short-lived EC2 test environments into that VPC without repeating network setup for each one, minimizing operational overhead.

https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpc-attachments.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html
https://docs.aws.amazon.com/vpc/latest/tgw/how-transit-gateways-work.html
https://docs.aws.amazon.com/en_en/vpc/latest/userguide/route-table-options.html

✗ **C** Create a new OU in AWS Organizations for testing. Create an AWS CloudFormation template that contains a VPC, necessary networking resources, a transit gateway attachment, and related routing configurations. Create a CloudFormation stack set that includes this template. Use CloudFormation StackSets for deployments into each account under the testing 01.1. Create a new account for each test environment.
Creating a separate account, VPC, and transit gateway attachment for each test environment introduces substantial setup and lifecycle management. StackSets can deploy resources across accounts, but that does not make this per-environment account and networking design lower overhead than sharing one VPC.

https://docs.aws.amazon.com/AWSCloudFormation/latest/TemplateReference/aws-properties-cloudformation-stackset-deploymenttargets.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html
https://docs.aws.amazon.com/he_il/AWSCloudFormation/latest/UserGuide/stacksets-orgs-associate-stackset-with-org.html

✗ **D** Convert the test environment EC2 instances into Docker images. Use AWS CloudFormation to configure an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in a new VPC, create a transit gateway attachment, and create related routing configurations. Use Kubernetes to manage the deployment and lifecycle of the test environments.
The requirement is to manage short-lived EC2-based environments, so converting them to containers and introducing an EKS cluster adds components and operational complexity that are not needed. The new VPC and transit gateway attachment also require networking setup, whereas a shared VPC can support the instances directly.

https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpc-attachments.html
https://docs.aws.amazon.com/eks/latest/userguide/eks-networking.html

Free demo · full exam: 681 questions

Question 35

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet reliability requirements.

A mobile gaming company is expanding into the global market. The company's game servers run in the us-east-1 Region. The game's client application uses UDP to communicate with the game servers and needs to be able to connect to a set of static IP addresses.

The company wants its game to be accessible on multiple continents. The company also wants the game to maintain its network performance and global availability.

Which solution meets these requirements?

A. Provision an Application Load Balancer (ALB) in front of the game servers. Create an Amazon CloudFront distribution that has no geographical restrictions. Set the ALB as the origin. Perform DNS lookups for the cloudfront.net domain name. Use the resulting IP addresses in the game's client application.

B. Provision game servers in each AWS Region. Provision an Application Load Balancer in front of the game servers. Create an Amazon Route 53 latency-based routing policy for the game's client application to use with DNS lookups.

C. Provision game servers in each AWS Region. Provision a Network Load Balancer (NLB) in front of the game servers. Create an accelerator in AWS Global Accelerator, and configure endpoint groups in each Region. Associate the NLBs with the corresponding Regional endpoint groups. Point the game client's application to the Global Accelerator endpoints.

D. Provision game servers in each AWS Region. Provision a Network Load Balancer (NLB) in front of the game servers. Create an Amazon CloudFront distribution that has no geographical restrictions. Set the NLB as the origin. Perform DNS lookups for the cloudfront.net domain name. Use the resulting IP addresses in the game's client application.

Answer: C. Provision game servers in each AWS Region. Provision a Network Load Balancer (NLB) in front of the game servers. Create an accelerator in AWS Global Accelerator, and configure endpoint groups in each Region. Associate the NLBs with the corresponding Regional endpoint groups. Point the game client's application to the Global Accelerator endpoints.

Explanation

Option C meets the requirements because AWS Global Accelerator provides static anycast IP addresses and directs client traffic to healthy endpoints in the appropriate Regions. Network Load Balancers support UDP, and Global Accelerator can use NLBs as regional endpoints. Deploying game servers across multiple Regions also provides the global reach and availability the company wants.

https://docs.aws.amazon.com/global-accelerator/latest/dg/about-endpoints.html
https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-listeners.html
https://docs.aws.amazon.com/us_en/global-accelerator/latest/dg/introduction-components.html
https://docs.aws.amazon.com/pdfs/global-accelerator/latest/dg/global-accelerator-guide.pdf

✗ **A** Provision an Application Load Balancer (ALB) in front of the game servers. Create an Amazon CloudFront distribution that has no geographical restrictions. Set the ALB as the origin. Perform DNS lookups for the cloudfront.net domain name. Use the resulting IP addresses in the game's client application.
This does not meet the UDP requirement: an ALB and CloudFront are not the suitable UDP forwarding path for this game traffic. CloudFront DNS results are also not the static IP addresses the client requires; Global Accelerator is designed to provide static anycast IP addresses.

https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/how-elastic-load-balancing-works.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/us_en/global-accelerator/latest/dg/introduction-components.html

✗ **B** Provision game servers in each AWS Region. Provision an Application Load Balancer in front of the game servers. Create an Amazon Route 53 latency-based routing policy for the game's client application to use with DNS lookups.
Regional game servers and latency-based DNS routing could direct clients toward a nearby Region, but the ALB does not satisfy the game's UDP communication requirement. This approach also does not provide the client with a set of static IP addresses as required.

https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/how-elastic-load-balancing-works.html
https://docs.aws.amazon.com/hi_in/Route53/latest/DeveloperGuide/routing-policy-latency.html

✓ **C** Provision game servers in each AWS Region. Provision a Network Load Balancer (NLB) in front of the game servers. Create an accelerator in AWS Global Accelerator, and configure endpoint groups in each Region. Associate the NLBs with the corresponding Regional endpoint groups. Point the game client's application to the Global Accelerator endpoints.
NLB supports UDP listeners, and Global Accelerator supports NLBs as regional endpoints. The accelerator provides static anycast IP addresses and routes traffic to regional endpoints, meeting the client's addressing, performance, and global availability requirements.

https://docs.aws.amazon.com/global-accelerator/latest/dg/about-endpoints.html
https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-listeners.html
https://docs.aws.amazon.com/us_en/global-accelerator/latest/dg/introduction-components.html
https://docs.aws.amazon.com/pdfs/global-accelerator/latest/dg/global-accelerator-guide.pdf

✗ **D** Provision game servers in each AWS Region. Provision a Network Load Balancer (NLB) in front of the game servers. Create an Amazon CloudFront distribution that has no geographical restrictions. Set the NLB as the origin. Perform DNS lookups for the cloudfront.net domain name. Use the resulting IP addresses in the game's client application.
Although NLBs can handle UDP, CloudFront is not the UDP acceleration solution for this game traffic. The IP addresses returned by CloudFront DNS are not the static anycast IP addresses needed by the client; Global Accelerator provides those addresses and can route to the regional NLBs.

https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-listeners.html
https://docs.aws.amazon.com/us_en/global-accelerator/latest/dg/introduction-components.html

Dumpsgate

Question 36

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet cost optimization goals.

A company needs to optimize the infrastructure for an application that uploads data to Amazon S3.

The uploads average 64 KB in size. When the data is uploaded, Amazon S3 sends an event to Amazon EventBridge. EventBridge then invokes an Amazon ECS application task.

The ECS task processes the data and stores the results in an Amazon DynamoDB table. Processing takes an average of 15 minutes. The company must keep the S3 data for 5 years and must keep the DynamoDB data for 15 days.

The application is gaining more users and is handling millions of S3 uploads every hour.

Which set of changes will provide the MOST cost-effective solution for the application?

A. Replace the ECS task with an AWS Lambda function for processing. Create S3 Lifecycle rules to move the S3 objects to S3 Intelligent-Tiering after 1 day and to expire the objects after 5 years. Configure DynamoDB Standard-Infrequent Access for the DynamoDB table.

B. Replace the S3 bucket with Amazon Managed Streaming for Apache Kafka (Amazon MSK) to receive the data. Configure tiered storage for data that is older than 1 day. Configure EventBridge to read messages from Amazon MSK in batches of 1,000 messages. Replace the ECS task with an AWS Lambda function for processing. Configure a TTL of 15 days on the DynamoDB table.

C. Create an Amazon Data Firehose stream to receive the data. Configure buffering to deliver messages every minute to Amazon S3 in gzip format. Purchase a Compute Savings Plan based on usage recommendations. Create S3 Lifecycle rules to move the S3 objects to S3 Glacier Deep Archive after 1 day and to expire the objects after 5 years. Configure a TTL of 15 days on the DynamoDB table.

D. Purchase a Compute Savings Plan based on usage recommendations. Create S3 Lifecycle rules to move the S3 objects to S3 Glacier Deep Archive after 1 day and to expire the objects after 5 years. Configure DynamoDB Standard-Infrequent Access for the DynamoDB table.

Answer: C. Create an Amazon Data Firehose stream to receive the data. Configure buffering to deliver messages every minute to Amazon S3 in gzip format. Purchase a Compute Savings Plan based on usage recommendations. Create S3 Lifecycle rules to move the S3 objects to S3 Glacier Deep Archive after 1 day and to expire the objects after 5 years. Configure a TTL of 15 days on the DynamoDB table.

Explanation

Option C reduces the overhead of handling millions of small uploads by buffering records and delivering compressed batches to S3. The lifecycle rules retain the data for the required five years while moving it to low-cost archival storage, and the DynamoDB TTL enforces the 15-day retention period. A Compute Savings Plan can also reduce the cost of the application's ongoing compute usage without moving its 15-minute processing task to a Lambda function.

https://docs.aws.amazon.com/firehose/latest/dev/buffering.html
https://docs.aws.amazon.com/pdfs/firehose/latest/dev/firehose-dg.pdf
https://docs.aws.amazon.com/AmazonS3/latest/userguide/how-to-set-lifecycle-configuration-intro.html

✗ **A** Replace the ECS task with an AWS Lambda function for processing. Create S3 Lifecycle rules to move the S3 objects to S3 Intelligent-Tiering after 1 day and to expire the objects after 5 years. Configure DynamoDB Standard-Infrequent Access for the DynamoDB table.
An average processing time of 15 minutes is at Lambda's maximum execution timeout, leaving no margin for slower executions or overhead. Intelligent-Tiering is not as cost-focused for five-year archival as

Free demo · full exam: 681 questions

moving the data to Glacier Deep Archive, and DynamoDB Standard-Infrequent Access does not itself expire items after 15 days.

https://docs.aws.amazon.com/lambda/latest/dg/configuration-timeout.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-transition-general-considerations.html
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.TableClasses.html

✗ **B** Replace the S3 bucket with Amazon Managed Streaming for Apache Kafka (Amazon MSK) to receive the data. Configure tiered storage for data that is older than 1 day. Configure EventBridge to read messages from Amazon MSK in batches of 1,000 messages. Replace the ECS task with an AWS Lambda function for processing. Configure a TTL of 15 days on the DynamoDB table.
Replacing S3 with MSK does not provide the specified S3-based five-year retention solution, and it introduces a streaming platform for data that must be retained as an archive. The replacement Lambda function also faces the 15-minute execution-time limit for processing that already averages 15 minutes. The DynamoDB TTL addresses its retention requirement, but does not offset these issues.

https://docs.aws.amazon.com/lambda/latest/dg/configuration-timeout.html

✓ **C** Create an Amazon Data Firehose stream to receive the data. Configure buffering to deliver messages every minute to Amazon S3 in gzip format. Purchase a Compute Savings Plan based on usage recommendations. Create S3 Lifecycle rules to move the S3 objects to S3 Glacier Deep Archive after 1 day and to expire the objects after 5 years. Configure a TTL of 15 days on the DynamoDB table.
Firehose buffering groups the high volume of small uploads into fewer S3 deliveries, and gzip reduces the amount of data stored. A lifecycle rule provides low-cost archival while retaining the data for five years, and DynamoDB TTL supports the 15-day retention requirement. A Compute Savings Plan can reduce ongoing compute costs while keeping the existing ECS processing task.

https://docs.aws.amazon.com/firehose/latest/dev/buffering.html
https://docs.aws.amazon.com/pdfs/firehose/latest/dev/firehose-dg.pdf
https://docs.aws.amazon.com/AmazonS3/latest/userguide/how-to-set-lifecycle-configuration-intro.html

✗ **D** Purchase a Compute Savings Plan based on usage recommendations. Create S3 Lifecycle rules to move the S3 objects to S3 Glacier Deep Archive after 1 day and to expire the objects after 5 years. Configure DynamoDB Standard-Infrequent Access for the DynamoDB table.
This option does not batch or compress the millions of small uploads, so it misses an opportunity to reduce S3 delivery and storage overhead. DynamoDB Standard-Infrequent Access is a table class, not a 15-day expiration mechanism, so the required DynamoDB retention is not enforced. A TTL is needed to expire the table data after 15 days.

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.TableClasses.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/how-to-set-lifecycle-configuration-intro.html

Question 37

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet reliability requirements.

A company hosts an ecommerce site using EC2, ALB, and DynamoDB in one AWS Region. The site uses a custom domain in Route 53. The company wants toreplicate the stack to a second Regionfordisaster recoveryandfaster accessfor global customers.

What should the architect do?

A. Use CloudFormation to deploy to the second Region. Use Route 53 latency-based routing. Enable global tables in DynamoDB.

B. Use the console to recreate the infra manually in the second Region. Use weighted routing.

C. Replicate only the S3 and DynamoDB data. Use Route 53 failover routing.

D. Use Beanstalk and DynamoDB Streams for replication. Use latency-based routing.

Answer: A. Use CloudFormation to deploy to the second Region. Use Route 53 latency-based routing. Enable global tables in DynamoDB.

Explanation

A CloudFormation deployment makes it possible to reproduce the application infrastructure in the second Region, including the EC2 and ALB components. Route 53 latency-based routing can direct customers to the Region expected to give them lower latency, supporting faster global access. DynamoDB global tables replicate table data across Regions, supporting the disaster-recovery requirement. [2] [3] [6] [9]

https://docs.aws.amazon.com/us_en/amazondynamodb/latest/developerguide/V2globaltables.tutorial.html
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-global-table-design.html
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables-CoreConcepts.html
https://docs.aws.amazon.com/pdfs/Route53/latest/DeveloperGuide/route53-dg.pdf

✓ **A** Use CloudFormation to deploy to the second Region. Use Route 53 latency-based routing. Enable global tables in DynamoDB.
This option replicates the application stack in the second Region and uses latency-based routing to direct customers to a lower-latency endpoint. DynamoDB global tables replicate table data across Regions, supporting the multi-Region disaster-recovery design. [2] [3] [6] [9]

https://docs.aws.amazon.com/us_en/amazondynamodb/latest/developerguide/V2globaltables.tutorial.html
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-global-table-design.html
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables-CoreConcepts.html
https://docs.aws.amazon.com/pdfs/Route53/latest/DeveloperGuide/route53-dg.pdf

✗ **B** Use the console to recreate the infra manually in the second Region. Use weighted routing.
Manual recreation is less repeatable than deploying the stack from a template, and weighted routing distributes traffic according to configured weights rather than selecting a Region based on customer latency. This does not address the stated goal of faster access for global customers as directly as latency-based routing. [1]

https://docs.aws.amazon.com/en_en/Route53/latest/DeveloperGuide/routing-policy-weighted.html

✗ **C** Replicate only the S3 and DynamoDB data. Use Route 53 failover routing.
The described stack also includes EC2 and an ALB, so replicating only data would leave key application infrastructure unavailable in the second Region. Failover routing is intended to direct traffic when an endpoint is unavailable, not to select a lower-latency Region for normal global access. [5]

https://docs.aws.amazon.com/en_en/Route53/latest/DeveloperGuide/dns-failover-types.html

✗ **D** Use Beanstalk and DynamoDB Streams for replication. Use latency-based routing.
This substitutes Elastic Beanstalk for the existing EC2-and-ALB stack and proposes Streams-based replication instead of the DynamoDB global tables suited to replicating table data across Regions. Although latency-based routing fits the global-access goal, the rest of this option does not provide the appropriate replication approach for the stated stack. [2] [3] [6] [9]

https://docs.aws.amazon.com/us_en/amazondynamodb/latest/developerguide/V2globaltables.tutorial.html
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-global-table-design.html
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables-CoreConcepts.html
https://docs.aws.amazon.com/pdfs/Route53/latest/DeveloperGuide/route53-dg.pdf

Free demo · full exam: 681 questions

Question 38

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet performance requirements.

A company runs a website on Amazon ECS containers that use the AWS Fargate launch type. The company configures AWS Application Auto Scaling by using a target tracking scaling policy. The company sets the request count as the scaling metric. An Application Load Balancer (ALB) serves traffic to the ECS containers. The website serves images on request and resizes the images to a predefined size to match the viewers' screens. After the website resizes an image, the website caches the image locally in a container and serves subsequent requests from the cache.

During periods of high traffic, the company observed that images load slowly and with high latency.

The company wants to minimize the latency to serve images.

Which solution will meet this requirement with the LEAST operational overhead?

A. Create a new Amazon CloudFront distribution and an Amazon S3 bucket. Set the ALB as one origin for the distribution and the S3 bucket as a second origin. Configure a cache behavior that routes image requests to the S3 origin, and configure a default cache behavior for the ALB origin. Pre-scale all images and upload the images to the S3 bucket.

B. Create an Amazon ElastiCache (Memcached) cluster. Update the application to read and write the resized images to the ElastiCache (Memcached) cluster by using the image name and size as the key.

C. Create an Amazon Aurora cluster and an Amazon S3 bucket. Update the application to store resized images in the S3 bucket and to store a cache key in the Aurora cluster. Configure the application to load the cache key from the Aurora cluster and to serve images from the S3 bucket.

D. Create an Amazon API Gateway HTTP API and enable API request caching. Replace the ALB with the HTTP API and remove the local caching in the application code.

Answer: A. Create a new Amazon CloudFront distribution and an Amazon S3 bucket. Set the ALB as one origin for the distribution and the S3 bucket as a second origin. Configure a cache behavior that routes image requests to the S3 origin, and configure a default cache behavior for the ALB origin. Pre-scale all images and upload the images to the S3 bucket.

Explanation

A CloudFront distribution can send image requests to an S3 origin while routing other requests to the ALB, so the application can continue serving dynamic website traffic. Pre-sized images in S3 can be cached at CloudFront edge locations, reducing the distance and work involved in serving repeated image requests. This avoids adding a separate cache cluster or database-backed cache and therefore meets the latency goal with less operational overhead.

https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_CacheBehavior.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistValuesCacheBehavior.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/RequestAndResponseBehaviorS3Origin.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistValuesOrigin.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/HowCloudFrontWorks.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/LocationsOfEdgeServers.html

✓ **A** Create a new Amazon CloudFront distribution and an Amazon S3 bucket. Set the ALB as one origin for the distribution and the S3 bucket as a second origin. Configure a cache behavior that routes image requests to the S3 origin, and configure a default cache behavior for the ALB origin. Pre-scale all images and upload the images to the S3 bucket.
CloudFront cache behaviors can route image requests to the S3 origin while sending other requests to the ALB, and CloudFront can serve cached content from edge locations closer to viewers. Pre-scaling

Free demo · full exam: 681 questions

images to the needed sizes lets the site serve them from S3 without repeatedly resizing them in ECS containers. This reduces image-serving latency without requiring the company to operate a separate cache cluster.

https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_CacheBehavior.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistValuesCacheBehavior.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/RequestAndResponseBehaviorS3Origin.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistValuesOrigin.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/HowCloudFrontWorks.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/LocationsOfEdgeServers.html

✗ **B** Create an Amazon ElastiCache (Memcached) cluster. Update the application to read and write the resized images to the ElastiCache (Memcached) cluster by using the image name and size as the key.
A Memcached cluster could hold resized images, but it requires application changes and operation of an additional cache tier. It also does not provide CloudFront's edge delivery to viewers, so it is not the least-overhead choice for minimizing image latency.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/LocationsOfEdgeServers.html

✗ **C** Create an Amazon Aurora cluster and an Amazon S3 bucket. Update the application to store resized images in the S3 bucket and to store a cache key in the Aurora cluster. Configure the application to load the cache key from the Aurora cluster and to serve images from the S3 bucket.
This design adds an Aurora database and application logic to look up cache keys, increasing operational overhead. It also does not put the image cache at CloudFront edge locations, so it does not address latency as directly as the S3-and-CloudFront solution.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/HowCloudFrontWorks.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/LocationsOfEdgeServers.html

✗ **D** Create an Amazon API Gateway HTTP API and enable API request caching. Replace the ALB with the HTTP API and remove the local caching in the application code.
API Gateway request caching would require replacing the existing ALB-based traffic path and changing the application's caching approach. It is not the edge-distribution solution provided by CloudFront for serving cached images closer to viewers, and the migration adds overhead rather than minimizing it.

https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/HowCloudFrontWorks.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/LocationsOfEdgeServers.html

Free demo · full exam: 681 questions

Question 39

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet cost optimization goals.

A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations lo manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.

Which solution is the MOST cost-effective way to meet these requirements?

A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in each account to create monthly reports for each business unit.

B. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.

C. Configure AWS Budgets in each account and configure budget alerts lhat are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.

D. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owner. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.

Explanation

Option B centralizes budget monitoring and reporting in the AWS Organizations management (master) account, which can view consolidated spending across the organization’s accounts. Cost Explorer there can provide monthly reports filtered using the existing application, environment, and owner tags. AWS Budgets can use those tags to set spending thresholds, and budget alerts can notify each business unit through its SNS topic. This avoids separately configuring the same reporting and alerting in every business unit account or building a custom reporting workflow.

https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-create-filters.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html
https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-sns-policy.html

✗ **A** Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in each account to create monthly reports for each business unit.
This distributes setup and reporting across the accounts instead of giving the governance team a centralized solution. Each account would need its own budgets and Cost Explorer reports, increasing

Free demo · full exam: 681 questions

administrative effort compared with managing consolidated spending from the organization's management account.

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html

✓ **B** Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
The management account can work with consolidated organization spending, so budgets and Cost Explorer reports can be managed centrally rather than duplicated in each account. Tag-based budget filters let the team monitor spending against thresholds, and SNS alerts can notify the appropriate business units. This meets both the monthly reporting and threshold-notification requirements with the least duplicated administration.

https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-create-filters.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html
https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-sns-policy.html

✗ **C** Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.
Like option A, this requires configuring budgets and reports separately in each account, so it does not provide the requested centralized solution. The account-by-account dashboard approach also duplicates work rather than using consolidated organization spending in the management account.

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-what-is.html

✗ **D** Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owner. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.
This introduces a custom Lambda-based processing and notification system even though AWS Budgets and Cost Explorer can meet the threshold-alert and reporting needs centrally. Building and maintaining that custom workflow is less cost-effective than using the managed billing tools and SNS notifications.

https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-create-filters.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html
https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-sns-policy.html

Question 40

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet performance requirements.

A company developed a pilot application by using AWS Elastic Beanstalk and Java. To save costs during development, the company's development team deployed the application into a singleinstance environment. Recent tests indicate that the application consumes more CPU than expected.

CPU utilization is regularly greater than 85%, which causes some performance bottlenecks.

A solutions architect must mitigate the performance issues before the company launches the application to production.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new Elastic Beanstalk application. Select a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the maximum CPU utilization is over 85% for 5 minutes.

B. Create a second Elastic Beanstalk environment. Apply the traffic-splitting deployment policy. Specify a percentage of incoming traffic to direct to the new environment in the average CPU utilization is over 85% for 5 minutes.

C. Modify the existing environment's capacity configuration to use a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.

D. Select the Rebuild environment action with the load balancing option Select an Availability Zones Add a scale-out rule that will run if the sum CPU utilization is over 85% for 5 minutes.

Answer: C. Modify the existing environment's capacity configuration to use a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.

Explanation

The existing single-instance environment can be changed to a load-balanced environment, allowing Elastic Beanstalk to run multiple instances and scale them based on CPU utilization. Configuring the average CPU utilization threshold to trigger scale-out addresses the reported sustained high CPU with less operational work than creating a separate application or environment. Selecting all Availability Zones provides the load-balanced environment with the specified availability-zone coverage.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/configuring-cluster-scaling.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environments-cfg-autoscaling-configuration-approaches.html

✗ **A** Create a new Elastic Beanstalk application. Select a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the maximum CPU utilization is over 85% for 5 minutes.
A new application is unnecessary because the existing environment can be changed to a load-balanced environment, so this adds avoidable setup and management. The rule also uses maximum CPU utilization, rather than the average CPU utilization specified in the correct approach.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/configuring-cluster-scaling.html

✗ **B** Create a second Elastic Beanstalk environment. Apply the traffic-splitting deployment policy. Specify a percentage of incoming traffic to direct to the new environment in the average CPU utilization is over 85%

Free demo · full exam: 681 questions

for 5 minutes.
Traffic splitting is a deployment policy for gradually directing traffic between application versions, not a means of automatically scaling capacity in response to CPU utilization. Creating and managing a second environment also adds overhead without addressing the underlying capacity configuration as directly as scaling the existing environment.

https://docs.aws.amazon.com/us_en/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/configuring-cluster-scaling.html

✓ **C** Modify the existing environment's capacity configuration to use a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.
Changing the existing environment to load-balanced lets it use multiple instances, and the scale-out rule can add capacity when average CPU utilization remains above the threshold. This directly addresses the CPU bottleneck while avoiding the extra setup of a new application or environment.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/configuring-cluster-scaling.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environments-cfg-autoscaling-configuration-approaches.html

✗ **D** Select the Rebuild environment action with the load balancing option Select an Availability Zones Add a scale-out rule that will run if the sum CPU utilization is over 85% for 5 minutes.
Rebuilding the environment is not required to change its capacity configuration, so it adds unnecessary operational work. The proposed sum CPU utilization trigger also does not match the average CPU utilization scale-out condition used to address the reported issue.

https://docs.aws.amazon.com/us_en/elasticbeanstalk/latest/dg/environment-management-rebuild.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environments-dashboard-actions.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/configuring-cluster-scaling.html

Question 41

Single choice · Domain: D2 Design for New Solutions › Determine a solution design to meet business requirements.

A startup company recently migrated a large ecommerce website to AWS The website has experienced a 70% increase in sates Software engineers are using a private GitHub repository to manage code The DevOps team is using Jenkins for builds and unit testing The engineers need to receive notifications for bad builds and zero downtime during deployments The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue The software engineers have decided to use AWS CodePipeline to manage their build and deployment process Which solution will meet these requirements'?

A. Use GitHub websockets to trigger the CodePipeline pipeline Use the Jenkins plugin for AWS CodeBuild to conduct unit testing Send alerts to an Amazon SNS topic for any bad builds Deploy inan in-place all-at-once deployment configuration using AWS CodeDeploy

B. Use GitHub webhooks to trigger the CodePipelme pipeline Use the Jenkins plugin for AWS CodeBuild to conduct unit testing Send alerts to an Amazon SNS topic for any bad builds Deploy in a blue'green deployment using AWS CodeDeploy

C. Use GitHub websockets to trigger the CodePipelme pipeline. Use AWS X-Ray for unit testing and static code analysis Send alerts to an Amazon SNS topic for any bad builds Deploy in a blue/green deployment using AWS CodeDeploy.

D. Use GitHub webhooks to trigger the CodePipeline pipeline Use AWS X-Ray for unit testing and static code analysis Send alerts to an Amazon SNS topic for any bad builds Deploy in an m-place. all-at-once deployment configuration using AWS CodeDeploy

Free demo · full exam: 681 questions

Answer: B. Use GitHub webhooks to trigger the CodePipelme pipeline Use the Jenkins plugin for AWS CodeBuild to conduct unit testing Send alerts to an Amazon SNS topic for any bad builds Deploy in a blue'green deployment using AWS CodeDeploy

Explanation

B meets the requirements by using GitHub webhooks to trigger CodePipeline and the Jenkins plugin for AWS CodeBuild to run unit tests. An SNS topic can notify the team about bad builds. A CodeDeploy blue/green deployment shifts users to the updated environment while keeping the existing environment available, supporting seamless deployment and rollback if a major issue occurs. [1][3][4][7]

https://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html
https://docs.aws.amazon.com/codebuild/latest/userguide/jenkins-plugin.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployments-rollback-and-redeploy.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/applications-create-blue-green.html

✗ **A** Use GitHub websockets to trigger the CodePipeline pipeline. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing. Send alerts to an Amazon SNS topic for any bad builds. Deploy in an in-place, all-at-once deployment configuration using AWS CodeDeploy.
The Jenkins plugin for AWS CodeBuild can support the unit-testing step, and SNS can be used for build notifications. However, this option uses websockets rather than GitHub webhooks to trigger the pipeline, and an in-place, all-at-once deployment does not provide the separate environment and traffic-switching approach needed for seamless blue/green deployment and rollback.

https://docs.aws.amazon.com/codedeploy/latest/userguide/deployments.html
https://docs.aws.amazon.com/codebuild/latest/userguide/jenkins-plugin.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployments-rollback-and-redeploy.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations-create.html

✓ **B** Use GitHub webhooks to trigger the CodePipeline pipeline. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing. Send alerts to an Amazon SNS topic for any bad builds. Deploy in a blue/green deployment using AWS CodeDeploy.
GitHub webhooks can trigger the pipeline when code changes, and the Jenkins plugin integrates Jenkins with CodeBuild for the testing step. SNS provides a channel for bad-build alerts. CodeDeploy blue/green deployment supports switching traffic to the new environment while retaining the previous environment for rollback if needed.

https://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html
https://docs.aws.amazon.com/codebuild/latest/userguide/jenkins-plugin.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployments-rollback-and-redeploy.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/applications-create-blue-green.html

✗ **C** Use GitHub websockets to trigger the CodePipeline pipeline. Use AWS X-Ray for unit testing and static code analysis. Send alerts to an Amazon SNS topic for any bad builds. Deploy in a blue/green deployment using AWS CodeDeploy.
Although blue/green deployment is appropriate for minimizing user disruption and supporting rollback, this option uses websockets instead of GitHub webhooks to trigger CodePipeline. AWS X-Ray is for tracing and analyzing application behavior, not for running unit tests or static code analysis.

https://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/applications-create-blue-green.html

✗ **D** Use GitHub webhooks to trigger the CodePipeline pipeline. Use AWS X-Ray for unit testing and static code analysis. Send alerts to an Amazon SNS topic for any bad builds. Deploy in an in-place, all-at-once deployment configuration using AWS CodeDeploy.
GitHub webhooks are suitable for triggering the pipeline, but AWS X-Ray does not conduct unit testing or static code analysis. Also, an in-place, all-at-once deployment is not the blue/green approach needed to keep the transition seamless and make rollback straightforward.

https://docs.aws.amazon.com/codedeploy/latest/userguide/deployments.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployments-rollback-and-redeploy.html
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations-create.html

Question 42

Single choice · Domain: D2 Design for New Solutions › Determine a solution design to meet business requirements.

A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.

Which solution will meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VPC. Create VPC peering connections that initiate from the central VPC to all other VPCs.

B. Create an AWS Direct Connect connection between the on-premises data center and AWS. Provision a transit VIF, and connect it to a Direct Connect gateway. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.

C. Create an AWS Site-to-Site VPN connection between the on-premises data centerand a new central VPC. Use a transit gateway with dynamic routing. Connect the transit gateway to all other VPCs.

D. Create an AWS Direct Connect connection between the on-premises data center and AWS Establish an AWS Site-to-Site VPN connection between all VPCs in each Region. Create VPC peering connections that initiate from the central VPC to all other VPCs.

Answer: B. Create an AWS Direct Connect connection between the on-premises data center and AWS. Provision a transit VIF, and connect it to a Direct Connect gateway. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.

Explanation

B meets the requirements by using Direct Connect for the on-premises connection and a transit VIF with a Direct Connect gateway to connect to transit gateways in multiple Regions. The transit gateways provide transitive routing between attached VPCs, unlike VPC peering. Direct Connect can provide higher, more consistent throughput than a VPN-based connection, and its data-transfer pricing can help reduce outbound network costs. [1][2][4][12]

https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html
https://docs.aws.amazon.com/en_en/directconnect/latest/UserGuide/direct-connect-transit-gateways.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/designing-control-tower-landing-zone/dc-transit-vif.html
https://docs.aws.amazon.com/whitepapers/latest/hybrid-connectivity/cost.html

✗ **A** Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VPC. Create VPC peering connections that initiate from the central VPC to all other VPCs.
VPC peering does not provide transitive routing through the central VPC, so this design cannot meet the requirement for routing between the VPC networks. It also uses a VPN rather than Direct Connect, so it does not address the throughput and consistent-network-experience requirements as well as the Direct Connect design. [3][5]

https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-basics.html
https://docs.aws.amazon.com/directconnect/latest/UserGuide/dc-ug.pdf

✓ **B** Create an AWS Direct Connect connection between the on-premises data center and AWS. Provision a transit VIF, and connect it to a Direct Connect gateway. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.
A transit VIF connects the Direct Connect connection to a Direct Connect gateway, which can associate with transit gateways in multiple Regions. Each regional transit gateway can connect the VPCs in that

Free demo · full exam: 681 questions

Region and provide transitive routing between them. Direct Connect also supports the scenario's goals for higher-throughput, more consistent connectivity and reduced data-transfer costs. [1][2][4][12]

https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html
https://docs.aws.amazon.com/en_en/directconnect/latest/UserGuide/direct-connect-transit-gateways.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/designing-control-tower-landing-zone/dc-transit-vif.html
https://docs.aws.amazon.com/whitepapers/latest/hybrid-connectivity/cost.html

✗ **C** Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VPC. Use a transit gateway with dynamic routing. Connect the transit gateway to all other VPCs.
A transit gateway can provide transitive routing among VPCs attached to it, but this option does not specify a multi-Region transit-gateway design for VPCs in different Regions. More importantly, it uses Site-to-Site VPN for the on-premises connection rather than Direct Connect, so it does not best meet the throughput, outbound-cost, and consistent-experience requirements. [2][5][12]

https://docs.aws.amazon.com/en_en/directconnect/latest/UserGuide/direct-connect-transit-gateways.html
https://docs.aws.amazon.com/directconnect/latest/UserGuide/dc-ug.pdf
https://docs.aws.amazon.com/whitepapers/latest/hybrid-connectivity/cost.html

✗ **D** Create an AWS Direct Connect connection between the on-premises data center and AWS Establish an AWS Site-to-Site VPN connection between all VPCs in each Region. Create VPC peering connections that initiate from the central VPC to all other VPCs.
This design does not establish the Direct Connect gateway and regional transit-gateway arrangement needed to connect the VPCs with transitive routing. VPC peering is nontransitive, and the additional VPN connections between VPCs do not make the central VPC's peering connections transitive. [2][3]

https://docs.aws.amazon.com/en_en/directconnect/latest/UserGuide/direct-connect-transit-gateways.html
https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-basics.html

Question 43

Single choice · Domain: D2 Design for New Solutions › Design a solution to meet performance requirements.

A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.

The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.

What could be the cause of the error messages for these customers?

A. The Lambda function reached its concurrency limit.
B. The Lambda function its Region limit for concurrency.
C. The company reached its API Gateway account limit for calls per second.
D. The company reached its API Gateway default per-method limit for calls per second.

Answer: C. The company reached its API Gateway account limit for calls per second.

Explanation

API Gateway enforces an account-level request-rate limit shared across APIs and methods in a Region.

Free demo · full exam: 681 questions

Premium API keys and their usage-plan capacity do not raise that account-level limit, so aggregate peak traffic can cause API Gateway to return 429 responses before it invokes Lambda. The fact that multiple methods are affected and the Lambda function is never invoked is consistent with throttling at API Gateway rather than a Lambda concurrency limit.

https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **A** The Lambda function reached its concurrency limit.
A Lambda concurrency limit can prevent invocations from being processed, but the logs indicate that Lambda is never invoked. The reported errors therefore point to throttling before the request reaches the function, at API Gateway.

https://docs.aws.amazon.com/lambda/latest/dg/monitoring-concurrency.html

✗ **B** The Lambda function its Region limit for concurrency.
A Region-level Lambda concurrency quota can restrict function invocations, but it would not explain requests being rejected before Lambda is invoked. The scenario's evidence instead points to API Gateway rejecting requests at its own throttling limit.

https://docs.aws.amazon.com/general/latest/gr/lambda-service.html
https://docs.aws.amazon.com/lambda/latest/dg/monitoring-concurrency.html

✓ **C** The company reached its API Gateway account limit for calls per second.
API Gateway applies an account-level request-rate limit shared by APIs and methods in a Region. Peak traffic can exhaust that limit even when individual customers have premium API keys, causing API Gateway to return 429 responses without invoking Lambda.

https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **D** The company reached its API Gateway default per-method limit for calls per second.
A per-method throttle applies to an individual method, whereas the reported errors affect multiple methods and are consistent with a shared account-level limit. API keys and usage-plan limits do not eliminate that broader account-level constraint.

https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html?trk=article-ssr-frontend-pulse_little-text-block

Dumpsgate

Question 44

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Improve the reliability of existing solutions.

A company runs an application in a primary AWS Region and maintains a pilot light disaster recovery (DR) environment in a secondary Region. The company uses an organization in AWS Organizations with all features enabled to manage its cloud environment. The company uses AWS CloudFormation to deploy resources to both Regions.

During a test, the company cannot fail over to the DR Region within the company's RTO target of 60 minutes because the DR Region's infrastructure has experienced configuration drift. The company observes drift in security groups, AMI versions, and application configurations. The configuration drift between the two Regions causes a delay of 30 to 45 minutes to scale up to production levels.

The company needs an automated solution to prevent configuration drift so the company can fail over to the DR Region within the target RTO.

Which solution will meet this requirement?

A. Deploy CloudFormation cross-Region stack sets in both Regions. Implement SCPs and IAM policies to restrict infrastructure modifications. Use Amazon EC2 Image Builder to create AMIs and distribute them to the DR Region. Configure AWS Config rules in both Regions. Set up Amazon EventBridge rules that invoke AWS Systems Manager Automation runbooks to remediate drift when drift is detected. Store application configurations in AWS Secrets Manager and enable cross-Region replication.

B. Use AWS Systems Manager Parameter Store to manage resource configurations centrally. Use AWS Lambda functions to synchronize AMIs between Regions daily. Use Amazon CloudWatch alarms to detect configuration changes. Use AWS CodePipeline to deploy configuration updates manually when drift is detected.

C. Deploy CloudFormation cross-Region stack sets in both Regions. Use SCPs to enforce configuration consistency. Use AWS Backup to replicate AMIs to the DR Region. Use AWS Config to scan for configuration drift daily. Use custom scripts to remediate configuration differences when drift is detected. Store application configurations in an Amazon S3 bucket and enable cross-Region replication to the DR Region.

D. Use Amazon EC2 Auto Scaling launch templates that have versioning enabled to maintain configuration consistency. Use AWS CloudTrail to log all configuration changes. Use AWS Health Dashboard to monitor infrastructure events. Use AWS Well-Architected Tool reviews quarterly to identify drift. Store application configurations in an Amazon S3 bucket and enable cross-Region replication to the DR Region.

Answer: A. Deploy CloudFormation cross-Region stack sets in both Regions. Implement SCPs and IAM policies to restrict infrastructure modifications. Use Amazon EC2 Image Builder to create AMIs and distribute them to the DR Region. Configure AWS Config rules in both Regions. Set up Amazon EventBridge rules that invoke AWS Systems Manager Automation runbooks to remediate drift when drift is detected. Store application configurations in AWS Secrets Manager and enable cross-Region replication.

Explanation

A combines preventive controls with automated detection and remediation so the DR Region stays aligned with the primary Region. CloudFormation StackSets can deploy the infrastructure across Regions,

Free demo · full exam: 681 questions

while SCPs and IAM policies restrict unauthorized changes. AWS Config can detect drift and EventBridge can trigger Systems Manager Automation to remediate it; EC2 Image Builder can distribute consistent AMIs, and Secrets Manager replication keeps application secrets available in the DR Region. Together, these measures reduce the configuration differences that caused the 30–45-minute failover delay.

https://docs.aws.amazon.com/secretsmanager/latest/userguide/replicate-secrets.html
https://docs.aws.amazon.com/imagebuilder/latest/userguide/how-image-builder-works.html
https://docs.aws.amazon.com/imagebuilder/latest/userguide/distribution-enhanced_functionality.html
https://docs.aws.amazon.com/imagebuilder/latest/userguide/manage-distribution-settings.html
https://docs.aws.amazon.com/imagebuilder/latest/userguide/cr-upd-ami-distribution-settings.html

✓ **A** Deploy CloudFormation cross-Region stack sets in both Regions. Implement SCPs and IAM policies to restrict infrastructure modifications. Use Amazon EC2 Image Builder to create AMIs and distribute them to the DR Region. Configure AWS Config rules in both Regions. Set up Amazon EventBridge rules that invoke AWS Systems Manager Automation runbooks to remediate drift when drift is detected. Store application configurations in AWS Secrets Manager and enable cross-Region replication.
This option addresses the observed sources of drift with consistent deployments, restrictions on unauthorized changes, and automated detection and remediation rather than waiting for a manual deployment. Image Builder can create and distribute standardized AMIs across Regions, and Secrets Manager supports cross-Region secret replication, helping the DR environment remain ready to scale up within the RTO.

https://docs.aws.amazon.com/secretsmanager/latest/userguide/replicate-secrets.html
https://docs.aws.amazon.com/imagebuilder/latest/userguide/how-image-builder-works.html
https://docs.aws.amazon.com/imagebuilder/latest/userguide/distribution-enhanced_functionality.html
https://docs.aws.amazon.com/imagebuilder/latest/userguide/manage-distribution-settings.html
https://docs.aws.amazon.com/imagebuilder/latest/userguide/cr-upd-ami-distribution-settings.html

✗ **B** Use AWS Systems Manager Parameter Store to manage resource configurations centrally. Use AWS Lambda functions to synchronize AMIs between Regions daily. Use Amazon CloudWatch alarms to detect configuration changes. Use AWS CodePipeline to deploy configuration updates manually when drift is detected.
Daily AMI synchronization can leave the DR Region using an outdated image between runs, and CloudWatch alarms alone do not remediate drift. The proposed manual CodePipeline deployment adds delay after drift is detected, which does not address the 30–45-minute delay that threatens the 60-minute RTO.

✗ **C** Deploy CloudFormation cross-Region stack sets in both Regions. Use SCPs to enforce configuration consistency. Use AWS Backup to replicate AMIs to the DR Region. Use AWS Config to scan for configuration drift daily. Use custom scripts to remediate configuration differences when drift is detected. Store application configurations in an Amazon S3 bucket and enable cross-Region replication to the DR Region.
Daily drift scans can allow configuration differences to persist until the next scan, and custom scripts do not provide the stated event-triggered remediation workflow in this option. AWS Backup and replicated S3 objects also do not provide the same purpose-built AMI distribution and managed secret replication described in A, so this is less suited to keeping the DR environment consistently ready for failover.

✗ **D** Use Amazon EC2 Auto Scaling launch templates that have versioning enabled to maintain configuration consistency. Use AWS CloudTrail to log all configuration changes. Use AWS Health Dashboard to monitor infrastructure events. Use AWS Well-Architected Tool reviews quarterly to identify drift. Store application configurations in an Amazon S3 bucket and enable cross-Region replication to the DR Region.
Versioned launch templates and CloudTrail can help track configurations and changes, but this option does not automatically detect and remediate drift across security groups, AMIs, and application configurations. Quarterly reviews are too infrequent to keep the DR environment within the 60-minute RTO, and S3 replication alone does not provide automated configuration remediation.

Free demo · full exam: 681 questions

Question 45

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a solution to improve security.

A company uses an organization in AWS Organizations to manage thousands of Amazon EC2 instances that run in multiple AWS accounts. The company groups the accounts into production and non-production environments that are in separate organizational units (OUs). The company identifies a critical vulnerability in its EC2 instances. The company needs to patch all affected instances. To comply with a security policy, the company must patch instances in the non-production environment before deploying the patch to the production environment. The company needs a scalable and auditable patch management solution across all accounts.

Which solution will meet these requirements with the LEAST operational overhead?

A. Tag each EC2 instance with an environment tag of either Prod or NonProd in all accounts. Enable Organizations delegated administration in the organization management account. Create an AWS Systems Manager Automation runbook that patches each instance based on the value of the environment tag. Use Amazon EventBridge to apply patching workflows for non-production instances first, followed by production instances.

B. Use AWS Systems Manager State Manager in each account to apply patches to EC2 instances during scheduled maintenance windows. Configure State Manager associations to non-production accounts first, followed by production accounts. Use AWS Config to produce compliance reports for all accounts.

C. Configure AWS Systems Manager Patch Manager with patch baselines. Enable delegated administration for Systems Manager and assign a delegated administrator account. Create Systems Manager Automation runbooks that target the OUs. Apply patches to instances in the non-production OU first by using approval gates before patching the production OU. Use Systems Manager Compliance and AWS CloudTrail to audit compliance in the OUs.

D. Use AWS Config rules across all accounts to manage patch compliance. Enable Organizations delegated administration in the organization management account. Create an AWS Lambda function to run a patching workflow in non-production accounts and then in production accounts. Configure the function to run when AWS Config identifies non-compliant accounts. Use AWS CloudTrail to audit all patch activity across the organization.

Answer: C. Configure AWS Systems Manager Patch Manager with patch baselines. Enable delegated administration for Systems Manager and assign a delegated administrator account. Create Systems Manager Automation runbooks that target the OUs. Apply patches to instances in the non- production OU first by using approval gates before patching the production OU. Use Systems Manager Compliance and AWS CloudTrail to audit compliance in the OUs.

Explanation

C uses Systems Manager Patch Manager and organization-level delegation to manage patching across accounts, while targeting the production and non-production OUs. Applying the patch to non-production first and requiring an approval gate before production enforces the required rollout order. Systems Manager Compliance and CloudTrail provide centralized compliance and activity records without requiring a custom patching workflow for each account.

https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-policies.html
https://docs.aws.amazon.com/systems-manager/latest/userguide/running-automations-multiple-accounts-regions.html?trk=article-ssr-frontend-pulse_little-text-block
https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/setting_up_delegated_admin.html
https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/patch-manager.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/pdfs/systems-manager/latest/userguide/systems-manager-ug.pdf

✗ **A** Tag each EC2 instance with an environment tag of either Prod or NonProd in all accounts. Enable Organizations delegated administration in the organization management account. Create an AWS Systems Manager Automation runbook that patches each instance based on the value of the environment tag. Use Amazon EventBridge to apply patching workflows for non-production instances first, followed by production instances.
This approach depends on maintaining environment tags on every instance and building custom runbooks and EventBridge workflows to sequence the rollout. The organization already groups accounts into OUs, so targeting those OUs with Systems Manager reduces this extra operational work.

https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/running-automations-map-targets.html
https://docs.aws.amazon.com/systems-manager/latest/userguide/running-automations-multiple-accounts-regions.html?trk=article-ssr-frontend-pulse_little-text-block
https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/setting_up_delegated_admin.html

✗ **B** Use AWS Systems Manager State Manager in each account to apply patches to EC2 instances during scheduled maintenance windows. Configure State Manager associations to non-production accounts first, followed by production accounts. Use AWS Config to produce compliance reports for all accounts.
Configuring State Manager associations in each account creates more account-by-account administration than using delegated Systems Manager management across the organization. The proposed sequencing of associations also does not provide the OU-targeted rollout and approval-gate approach described for the required controlled production deployment.

https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-policies.html
https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/setting_up_delegated_admin.html
https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/patch-manager.html

✓ **C** Configure AWS Systems Manager Patch Manager with patch baselines. Enable delegated administration for Systems Manager and assign a delegated administrator account. Create Systems Manager Automation runbooks that target the OUs. Apply patches to instances in the non- production OU first by using approval gates before patching the production OU. Use Systems Manager Compliance and AWS CloudTrail to audit compliance in the OUs.
Patch Manager provides the patching and baseline capabilities, and delegated administration supports managing Systems Manager across the organization. Targeting the OUs lets the company patch non-production first and hold the production rollout behind an approval gate, while Systems Manager Compliance and CloudTrail support auditing.

https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-policies.html
https://docs.aws.amazon.com/systems-manager/latest/userguide/running-automations-multiple-accounts-regions.html?trk=article-ssr-frontend-pulse_little-text-block
https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/setting_up_delegated_admin.html
https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/patch-manager.html
https://docs.aws.amazon.com/pdfs/systems-manager/latest/userguide/systems-manager-ug.pdf

✗ **D** Use AWS Config rules across all accounts to manage patch compliance. Enable Organizations delegated administration in the organization management account. Create an AWS Lambda function to run a patching workflow in non-production accounts and then in production accounts. Configure the function to run when AWS Config identifies non-compliant accounts. Use AWS CloudTrail to audit all patch activity across the organization.
This solution requires a custom Lambda patching workflow and event integration, increasing operational overhead. AWS Config can help identify compliance issues, but the option does not use Patch Manager's patching capabilities and organization-wide management approach for the rollout.

https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-policies.html
https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/setting_up_delegated_admin.html
https://docs.aws.amazon.com/en_en/systems-manager/latest/userguide/patch-manager.html

Question 46

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a solution to improve security.

A company is implementing a CI/CD pipeline for containerized applications by using Amazon ECR.

The company needs a solution that provides automated vulnerability scanning of all container images and the ability to prioritize remediation based on actual deployment status. The solution must scan repositories for operating system and programming language vulnerabilities. The solution must provide continuous monitoring as new Common Vulnerabilities and Exposures (CVEs) are discovered and integrate the findings into the company's existing event-driven workflow.

Which solution will meet these requirements with the LEAST operational overhead?

A. Turn on Amazon ECR basic scanning. Enable scan on push for all repositories. Configure Amazon EventBridge rules to capture scan completion events and route the findings to AWS Security Hub CSPM. Use ECR lifecycle policies to automatically delete images that have critical vulnerabilities.

B. Create AWS Lambda functions that are invoked by ECR image push events. Configure the Lambda functions to pull images and scan the images by using open-source tools. Store the scan results in Amazon DynamoDB. Create Amazon CloudWatch dashboards to visualize vulnerability trends across repositories.

C. Turn on Amazon ECR basic scanning and enhanced scanning. Configure manual scan triggers for basic scanning and enhanced scanning. Use AWS Systems Manager Parameter Store to maintain a list of approved base images. Create AWS Config rules to check that only approved images are deployed to production environments and automatically remediate non-compliant resources.

D. Turn on Amazon ECR enhanced scanning. Integrate enhanced scanning with Amazon Inspector. Configure continuous scanning for repositories. Use Amazon Inspector to track the deployment status of container images. Configure Amazon EventBridge rules to capture Amazon Inspector findings and route the findings to the company's existing event-driven workflow.

Answer: D. Turn on Amazon ECR enhanced scanning. Integrate enhanced scanning with Amazon Inspector. Configure continuous scanning for repositories. Use Amazon Inspector to track the deployment status of container images. Configure Amazon EventBridge rules to capture Amazon Inspector findings and route the findings to the company's existing event-driven workflow.

Explanation

Amazon ECR enhanced scanning integrates with Amazon Inspector to scan container images for operating system and programming language vulnerabilities. Continuous scanning keeps findings current as new CVEs are discovered, and Inspector can help prioritize findings using whether images are in use in supported container environments. EventBridge can route Inspector findings into the company's existing event-driven workflow, avoiding the operational burden of building and maintaining a custom scanner.

https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning-enhanced.html
https://docs.aws.amazon.com/inspector/latest/user/supported.html
https://docs.aws.amazon.com/inspector/latest/user/eventbridge-integration.html
https://docs.aws.amazon.com/inspector/latest/user/ecr-integration.html
https://docs.aws.amazon.com/inspector/latest/user/scanning-ecr.html

✗ **A** Turn on Amazon ECR basic scanning. Enable scan on push for all repositories. Configure Amazon EventBridge rules to capture scan completion events and route the findings to AWS Security Hub CSPM. Use ECR lifecycle policies to automatically delete images that have critical vulnerabilities.
Basic scanning with scan on push does not provide the continuous monitoring for newly discovered CVEs

Free demo · full exam: 681 questions

required here, and it does not provide the needed prioritization based on actual deployment status. Automatically deleting images based on severity also does not solve that prioritization requirement. Enhanced scanning with Inspector supports the required scanning and deployment-aware findings.

https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning-enhanced.html
https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **B** Create AWS Lambda functions that are invoked by ECR image push events. Configure the Lambda functions to pull images and scan the images by using open-source tools. Store the scan results in Amazon DynamoDB. Create Amazon CloudWatch dashboards to visualize vulnerability trends across repositories.
This approach requires building and operating custom scanning, storage, and visualization components, creating more operational overhead than using the managed ECR and Inspector integration. An image-push trigger also does not by itself provide continuous monitoring as new CVEs are discovered or prioritization based on deployment status.

https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning-enhanced.html
https://docs.aws.amazon.com/inspector/latest/user/ecr-integration.html
https://docs.aws.amazon.com/inspector/latest/user/scanning-ecr.html

✗ **C** Turn on Amazon ECR basic scanning and enhanced scanning. Configure manual scan triggers for basic scanning and enhanced scanning. Use AWS Systems Manager Parameter Store to maintain a list of approved base images. Create AWS Config rules to check that only approved images are deployed to production environments and automatically remediate non-compliant resources.
Manual scan triggers do not meet the requirement for continuous monitoring as new CVEs are discovered. Approved-image checks and remediation address deployment policy, but they do not provide the requested vulnerability prioritization based on actual deployment status or integrate Inspector findings into the existing event-driven workflow.

https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning-enhanced.html
https://docs.aws.amazon.com/inspector/latest/user/ecr-integration.html
https://docs.aws.amazon.com/inspector/latest/user/scanning-ecr.html

✓ **D** Turn on Amazon ECR enhanced scanning. Integrate enhanced scanning with Amazon Inspector. Configure continuous scanning for repositories. Use Amazon Inspector to track the deployment status of container images. Configure Amazon EventBridge rules to capture Amazon Inspector findings and route the findings to the company's existing event-driven workflow.
Enhanced scanning with Inspector covers operating system and programming language vulnerabilities and can continuously update findings as new CVEs are identified. Inspector provides image-use context that helps prioritize remediation based on deployment status, while EventBridge routes findings into existing event-driven workflows. This uses managed AWS capabilities rather than requiring custom scanning infrastructure.

https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning-enhanced.html
https://docs.aws.amazon.com/inspector/latest/user/supported.html
https://docs.aws.amazon.com/inspector/latest/user/eventbridge-integration.html
https://docs.aws.amazon.com/inspector/latest/user/ecr-integration.html
https://docs.aws.amazon.com/inspector/latest/user/scanning-ecr.html

Free demo · full exam: 681 questions

Question 47

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine how to improve the performance of existing solutions.

A company operates a data analytics platform that uses Amazon S3 to store petabytes of customer dat a. The platform stores data in multiple S3 buckets across four AWS Regions. The company has configured S3 Cross-Region Replication and uses several S3 prefixes across the S3 buckets.

The company wants to analyze and visualize trends and replication metrics for its S3 storage, including for specific prefixes, over a 12-month period.

Which solution will meet these requirements with the LEAST operational overhead?

A. Use the S3 Intelligent-Tiering storage class for all the S3 buckets. Configure Amazon CloudWatch to collect metrics for storage transitions. Track S3 API calls in AWS CloudTrail. Export all logs and metrics to an S3 bucket. Use Amazon Athena to query the logs. Add bucket prefixes to query clauses for the specified prefixes. Use Amazon QuickSight to generate visualizations.

B. Create an S3 Storage Lens dashboard. Enable advanced metrics and recommendations. Create Storage Lens groups and configure filters for the specified prefixes. Attach the groups to the dashboard. Configure Storage Lens metrics to export to an S3 bucket. Use Amazon Athena to query the metrics. Use Amazon QuickSight to generate visualizations for custom analytics.

C. Configure AWS Cost Explorer with S3 storage filters. Set up scheduled reports to produce daily exports. Export all data to Amazon CloudWatch Logs and use CloudWatch Logs Insights to query and visualize the data. Track the CloudWatch bucket prefix metric for the specified prefixes.

D. Create an S3 Storage Lens dashboard. Enable default metrics and recommendations. Configure S3 server access logging on all the buckets. Deliver all logs and metrics to an Amazon Managed Service for Prometheus workspace. Add bucket prefix metadata. Use Amazon Managed Grafana to query and visualize the data.

Answer: B. Create an S3 Storage Lens dashboard. Enable advanced metrics and recommendations. Create Storage Lens groups and configure filters for the specified prefixes. Attach the groups to the dashboard. Configure Storage Lens metrics to export to an S3 bucket. Use Amazon Athena to query the metrics. Use Amazon QuickSight to generate visualizations for custom analytics.

Explanation

B uses S3 Storage Lens to analyze storage trends across the company's buckets and Regions, while advanced metrics and recommendations provide additional metrics, including replication metrics. Storage Lens groups can filter metrics by prefixes, and exporting the metrics to S3 lets the company query them with Athena and visualize custom analyses in QuickSight. Storage Lens retains historical metrics for up to 15 months, which covers the requested 12-month period. This managed approach avoids building a separate pipeline from storage-transition metrics and API logs.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_view_metrics.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-lens-groups.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_creating_dashboard.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_understanding_metrics_export_schema.html

✗ **A** Use the S3 Intelligent-Tiering storage class for all the S3 buckets. Configure Amazon CloudWatch to collect metrics for storage transitions. Track S3 API calls in AWS CloudTrail. Export all logs and metrics to an S3 bucket. Use Amazon Athena to query the logs. Add bucket prefixes to query clauses for the specified prefixes. Use Amazon QuickSight to generate visualizations.

Free demo · full exam: 681 questions

This collects storage-transition metrics and API-call logs, but it does not provide the purpose-built Storage Lens view of storage and replication metrics across the buckets. Querying logs by prefix would also require the company to assemble and maintain its own analytics workflow rather than use Storage Lens groups. Intelligent-Tiering is not needed to analyze the requested trends.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_view_metrics.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-lens-groups.html

✓ **B** Create an S3 Storage Lens dashboard. Enable advanced metrics and recommendations. Create Storage Lens groups and configure filters for the specified prefixes. Attach the groups to the dashboard. Configure Storage Lens metrics to export to an S3 bucket. Use Amazon Athena to query the metrics. Use Amazon QuickSight to generate visualizations for custom analytics.
Storage Lens provides storage and replication metrics across S3, and advanced metrics provide the additional detail needed for this analysis. Storage Lens groups can filter metrics for the specified prefixes, and exporting metrics to S3 enables queries with Athena and custom visualizations with QuickSight. Its historical metrics retention covers the requested 12-month analysis period.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_view_metrics.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-lens-groups.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_creating_dashboard.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_understanding_metrics_export_schema.html

✗ **C** Configure AWS Cost Explorer with S3 storage filters. Set up scheduled reports to produce daily exports. Export all data to Amazon CloudWatch Logs and use CloudWatch Logs Insights to query and visualize the data. Track the CloudWatch bucket prefix metric for the specified prefixes.
Cost Explorer is focused on cost and usage analysis, not the S3 replication metrics requested here. The proposed CloudWatch bucket-prefix metric does not provide the described per-prefix replication analysis, and this approach adds a reporting and logging pipeline instead of using Storage Lens.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_view_metrics.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-lens-cloudwatch-metrics-dimensions.html

✗ **D** Create an S3 Storage Lens dashboard. Enable default metrics and recommendations. Configure S3 server access logging on all the buckets. Deliver all logs and metrics to an Amazon Managed Service for Prometheus workspace. Add bucket prefix metadata. Use Amazon Managed Grafana to query and visualize the data.
Default Storage Lens metrics do not provide the advanced replication detail needed for this analysis, and server access logs record requests rather than serving as a source of Storage Lens replication metrics. The proposed Prometheus delivery path also does not substitute for exporting Storage Lens metrics to S3 for Athena-based custom analysis. This adds operational components without meeting the stated requirements as directly as B.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_view_metrics.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerLogs.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html

Question 48

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a solution to improve security.

A healthcare company is building a user support chat-based assistant on Amazon Bedrock. Users will ask health questions that might include personal details in the prompts.

A solutions architect must configure a solution that can do the following:

- Prevent the assistant from providing medical diagnosis advice.
- Redact personally identifiable information (PII) from both user inputs and model responses.
- Enforce the same controls even if the company changes foundation models (FMs) later.
- Evaluate risky user prompts before sending the prompts to a model to avoid unnecessary inference costs.

Which solution will meet these requirements?

A. Store approved health support guidelines in an Amazon Bedrock knowledge base. Configure system prompts that instruct the model not to provide diagnosis advice. Use an AWS Lambda function after inference to remove PII from the model response before the response is returned to users.

B. Fine-tune an FM on approved support conversations. Add prompt templates that prohibit diagnosis advice. Run a separate review process that scans transcripts for prohibited topics and sensitive information after the conversations end.

C. Build a custom moderation layer in the application to inspect prompts for prohibited topics and to redact sensitive information from user inputs. Invoke the model through the Converse API. Use separate post-processing logic to redact sensitive information from responses before displaying responses to users.

D. Create an Amazon Bedrock guardrail. Configure denied topics for medical diagnosis advice. Configure sensitive information filters to mask PII. Configure content filters. Call the ApplyGuardrail API on user prompts before inference. Include the same guardrail in the Converse API to evaluate model responses.

Answer: D. Create an Amazon Bedrock guardrail. Configure denied topics for medical diagnosis advice. Configure sensitive information filters to mask PII. Configure content filters. Call the ApplyGuardrail API on user prompts before inference. Include the same guardrail in the Converse API to evaluate model responses.

Explanation

Amazon Bedrock Guardrails can apply denied-topic controls to prevent medical diagnosis advice and sensitive-information filters to mask PII. Calling ApplyGuardrail on each prompt lets the company evaluate it before model inference, avoiding inference costs for prompts that should be blocked. The same guardrail can be applied through Converse to evaluate model responses, keeping controls consistent even if the foundation model changes. References: 1, 2, 4, 5, 6.

https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-use-converse-api.html
https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-denied-topics.html
https://docs.aws.amazon.com/en_en/bedrock/latest/userguide/guardrails-use-independent-api.html
https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-sensitive-filters.html
https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-how.html

✗ **A** Store approved health support guidelines in an Amazon Bedrock knowledge base. Configure system prompts that instruct the model not to provide diagnosis advice. Use an AWS Lambda function after

Free demo · full exam: 681 questions

inference to remove PII from the model response before the response is returned to users.
A knowledge base and system prompt do not provide the same enforceable guardrail controls as a denied-topic policy, and the Lambda function only removes PII from responses, not user inputs. The checks also occur after inference, so risky prompts can still incur model costs. References: 2, 5, 6.

https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-denied-topics.html
https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-sensitive-filters.html
https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-how.html

✗ **B** Fine-tune an FM on approved support conversations. Add prompt templates that prohibit diagnosis advice. Run a separate review process that scans transcripts for prohibited topics and sensitive information after the conversations end.
Fine-tuning and prompt templates do not reliably enforce the requested controls, and reviewing transcripts after conversations end is too late to block risky prompts before inference. This approach also does not redact PII from both inputs and responses in real time. Reference: 3.

https://docs.aws.amazon.com/bedrock/latest/userguide/reinforcement-fine-tuning.html

✗ **C** Build a custom moderation layer in the application to inspect prompts for prohibited topics and to redact sensitive information from user inputs. Invoke the model through the Converse API. Use separate post-processing logic to redact sensitive information from responses before displaying responses to users.
The custom layer can inspect and redact inputs before inference, but the stated response post-processing only redacts sensitive information; it does not specify blocking diagnosis advice in model responses. Separate custom logic also does not provide the shared Bedrock guardrail applied to both prompt checks and Converse responses. References: 1, 4, 5.

https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-use-converse-api.html
https://docs.aws.amazon.com/en_en/bedrock/latest/userguide/guardrails-use-independent-api.html
https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-sensitive-filters.html

✓ **D** Create an Amazon Bedrock guardrail. Configure denied topics for medical diagnosis advice. Configure sensitive information filters to mask PII. Configure content filters. Call the ApplyGuardrail API on user prompts before inference. Include the same guardrail in the Converse API to evaluate model responses.
A denied topic targets medical diagnosis advice, and sensitive-information filters can mask PII. ApplyGuardrail evaluates prompts independently before inference, while using that same guardrail with Converse evaluates responses too. This centralizes the controls rather than tying them to a particular foundation model. References: 1, 2, 4, 5, 6.

https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-use-converse-api.html
https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-denied-topics.html
https://docs.aws.amazon.com/en_en/bedrock/latest/userguide/guardrails-use-independent-api.html
https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-sensitive-filters.html
https://docs.aws.amazon.com/bedrock/latest/userguide/guardrails-how.html

Free demo · full exam: 681 questions

Question 49

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Improve the reliability of existing solutions.

A company runs a video-on-demand (VOD) content streaming application on AWS. The application includes an Amazon CloudFront distribution that uses the default cache behavior. The distribution has a single origin that points to an Amazon S3 bucket that contains the video files.

The company wants to improve the application's reliability. The company creates a second S3 bucket and configures S3 Cross-Region Replication (CRR) between the S3 buckets. The company must implement high availability for the CloudFront deployment and must ensure that failover begins within 1 second.

Which change to the current architecture will meet these requirements with the LEAST operational overhead?

A. Create a second CloudFront distribution that uses the second S3 bucket as a single origin. Create an origin group. Add both distributions to the origin group. Set the original distribution as the primary distribution. Set the new distribution as the secondary distribution. Create an Amazon Route 53 health check to monitor the health of the primary distribution and secondary distribution every second.

B. Create a new origin in the existing CloudFront distribution. Specify the second S3 bucket as the new origin. Create an origin group. Add the original origin as the primary origin. Add the new origin as the secondary origin. Set the origin response timeout value to 1. Set the origin connection attempts value to 1.

C. Create a new origin in the existing CloudFront distribution. Specify the second S3 bucket as the new origin. Create an origin group. Add the original origin as the primary origin. Add the new origin as the secondary origin. Update the default cache behavior to use the origin group. Set the origin connection timeout value to 1. Set the origin connection attempts value to 1.

D. Create a new origin in the existing CloudFront distribution. Specify the second S3 bucket as the new origin. Create an AWS Lambda function to monitor the health of the original origin. Program the Lambda function to update the CloudFront distribution and promote the secondary origin to primary if a health check fails. Create an Amazon EventBridge scheduled rule to invoke the Lambda function every second.

Answer: C. Create a new origin in the existing CloudFront distribution. Specify the second S3 bucket as the new origin. Create an origin group. Add the original origin as the primary origin. Add the new origin as the secondary origin. Update the default cache behavior to use the origin group. Set the origin connection timeout value to 1. Set the origin connection attempts value to 1.

Explanation

A CloudFront origin group provides failover between the two S3 origins within the existing distribution, with the original bucket primary and the replicated bucket secondary. The default cache behavior must use that origin group for viewer requests to receive the failover behavior. Setting the origin connection timeout and connection attempts to 1 minimizes the time CloudFront waits before trying the secondary origin, while avoiding separate monitoring or distribution-management automation.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistValuesOrigin.html
https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_OriginGroup.html

✗ **A** Create a second CloudFront distribution for the second bucket, add both distributions to an origin

Free demo · full exam: 681 questions

group, and use Route 53 health checks every second.
CloudFront origin groups contain origins, not separate CloudFront distributions, so this does not configure the requested native origin failover. Route 53 health checks also add operational components and do not provide the stated every-second failover mechanism.

https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_OriginGroup.html
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-determining-health-of-endpoints.html

✗ **B** Add the second S3 bucket as an origin and create an origin group, but do not update the default cache behavior; set the origin response timeout and connection attempts to 1.
The default cache behavior must use the origin group; otherwise requests using that behavior will not be routed through the group for failover. Also, the response timeout is not the connection timeout setting specified in the correct configuration.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistValuesOrigin.html

✓ **C** Add the second S3 bucket as an origin, create an origin group with the original origin primary and the new origin secondary, update the default cache behavior to use the group, and set the origin connection timeout and connection attempts to 1.
This uses CloudFront's native origin failover with the replicated bucket as the secondary origin, and attaches the group to the behavior serving the video requests. Setting the connection timeout and attempts to 1 minimizes the wait before CloudFront tries the secondary origin, without requiring separate health-check or failover automation.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistValuesOrigin.html
https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_OriginGroup.html

✗ **D** Use a Lambda function to monitor the original origin and update the CloudFront distribution to promote the secondary origin, invoked by an EventBridge rule every second.
This adds monitoring, Lambda, scheduled invocation, and CloudFront distribution updates, creating substantially more operational overhead than CloudFront origin groups. It also relies on changing distribution configuration after a failure rather than using the native origin failover mechanism for requests.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html

Question 50

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a strategy to improve overall operational excellence.

A retailer is deploying a customer service AI agent that uses multiple tools to look up orders, process refunds, and make address changes. The retailer needs a solution for its CI/CD pipeline that blocks promotion if a new release regresses on specific metrics. The specific metrics include task completion, tool selection, and final response quality. The retailer wants to test the metric evaluation against a set of golden interaction traces. After deployment, the solution must also continuously monitor for agent degradation and perform a periodic human audit of production results.

Which solution will meet these requirements with the LEAST operational overhead?

A. Build custom checks by using Amazon CloudWatch Logs and AWS Lambda to parse tool calls and response text. Use weekly manual reviews to decide whether the deployed AI agent meets performance standards.

B. Use an Amazon Bedrock model evaluation job on the underlying foundation model (FM) by using prompt-response pairs. Promote the release if the model's helpfulness and correctness scores improve over the previous release.

C. Configure agent trace logging. Add Amazon Bedrock AgentCore Evaluations to the pipeline for on-demand evaluation by using built-in and custom evaluators. Enable online evaluation after deployment. Periodically review a sampled subset of sessions.

D. Use Amazon Bedrock AgentCore Evaluations in online mode after deployment by using built-in evaluators. Use post-deployment rollback if the evaluation detects regression.

Answer: C. Configure agent trace logging. Add Amazon Bedrock AgentCore Evaluations to the pipeline for on- demand evaluation by using built-in and custom evaluators. Enable online evaluation after deployment. Periodically review a sampled subset of sessions.

Explanation

C is the best fit because AgentCore Evaluations can run on demand in the CI/CD pipeline against golden interaction traces, using built-in and custom evaluators for the agent's task completion, tool selection, and final response quality. After deployment, online evaluations can continuously monitor production interactions, while periodic reviews of sampled sessions provide the requested human audit. This uses managed evaluation capabilities rather than building custom checks and manual release decisions.

https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/custom-evaluators.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/online-evaluations.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/evaluators.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/on-demand-evaluations.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/evaluations.html

✗ **A** Build custom checks by using Amazon CloudWatch Logs and AWS Lambda to parse tool calls and response text. Use weekly manual reviews to decide whether the deployed AI agent meets performance standards.
This requires custom parsing and evaluation logic, increasing operational overhead. Weekly manual reviews also do not provide the requested automated, pre-promotion evaluation against golden traces or continuous post-deployment monitoring.

✗ **B** Use an Amazon Bedrock model evaluation job on the underlying foundation model (FM) by using prompt-response pairs. Promote the release if the model's helpfulness and correctness scores improve over the previous release.
This evaluates the underlying model with prompt-response pairs and helpfulness or correctness metrics,

Free demo · full exam: 681 questions

rather than evaluating the agent's tool selection and end-to-end task completion on interaction traces. It also does not provide the requested continuous monitoring and periodic human audit of production results.

https://docs.aws.amazon.com/bedrock/latest/userguide/model-evaluation-metrics.html
https://docs.aws.amazon.com/bedrock/latest/userguide/model-evaluation-prompt-datasets-judge.html
https://docs.aws.amazon.com/bedrock/latest/userguide/model-evaluation-type-judge-prompt.html

✓ **C** Configure agent trace logging. Add Amazon Bedrock AgentCore Evaluations to the pipeline for on-demand evaluation by using built-in and custom evaluators. Enable online evaluation after deployment. Periodically review a sampled subset of sessions.
On-demand AgentCore Evaluations can assess golden interaction traces in the pipeline, and custom evaluators can cover agent-specific metrics such as tool selection and task completion alongside built-in evaluations. Online evaluations support ongoing production monitoring, and reviewing sampled sessions provides the periodic human audit the retailer requires.

https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/custom-evaluators.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/online-evaluations.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/evaluators.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/on-demand-evaluations.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/create-online-evaluations.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/evaluations.html

✗ **D** Use Amazon Bedrock AgentCore Evaluations in online mode after deployment by using built-in evaluators. Use post-deployment rollback if the evaluation detects regression.
Online evaluation can monitor the agent after deployment, but this option does not evaluate a release against golden traces before promotion in the CI/CD pipeline. It also specifies only built-in evaluators, which might not cover the required agent-specific metrics, and it omits the periodic human audit.

https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/custom-evaluators.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/online-evaluations.html
https://docs.aws.amazon.com/bedrock-agentcore/latest/devguide/on-demand-evaluations.html

Question 51

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine how to improve the performance of existing solutions.

A company uses Amazon CloudFront to serve a static application from an Amazon S3 bucket origin. A custom domain maps to the CloudFront distribution. The application calls an Amazon API Gateway Regional API endpoint in the ap-northeast-1 Region for backend APIs.

The company configures cross-origin resource sharing (CORS) so the application can access the APIs.

Tests find that API calls are slower from outside ap-northeast-1, which causes slow performance.

The company must improve the performance of the application.

Which solution will meet this requirement?

A. Enable CloudFront Origin Shield and S3 Transfer Acceleration for the S3 bucket.

B. Use an Amazon Route 53 latency-based routing policy for the Regional API endpoint and the CloudFront distribution.

C. Set the Regional API endpoint as an origin for the CloudFront distribution. Access the backend APIs by using the application's DNS name.

D. Set up AWS Global Accelerator in the same account that hosts the CloudFront distribution. Access the backend APIs by using Global Accelerator.

Free demo · full exam: 681 questions

Answer: C. Set the Regional API endpoint as an origin for the CloudFront distribution. Access the backend APIs by using the application's DNS name.

Explanation

Configure the Regional API endpoint as an origin in the existing CloudFront distribution, and route API requests through the application's DNS name using an appropriate cache behavior. This lets requests use CloudFront's edge network on the way to the API Gateway endpoint in ap-northeast-1, improving performance for users farther from that Region. The existing CORS configuration still needs to allow the application's requests; this solution changes the path used to reach the API, not the API's Region. [1][2][3] [12]

https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_Origin.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/RequestAndResponseBehaviorCustomOrigin.html
https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_DefaultCacheBehavior.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistValuesOrigin.html

✗ **A** Enable CloudFront Origin Shield and S3 Transfer Acceleration for the S3 bucket.
These features address delivery or origin access for the S3-hosted static application, not the API Gateway calls that are slow. The API calls would still go directly to the Regional API endpoint in ap-northeast-1.

✗ **B** Use an Amazon Route 53 latency-based routing policy for the Regional API endpoint and the CloudFront distribution.
A latency-based routing policy does not make the single Regional API endpoint available closer to users or route API requests through CloudFront. The API calls would still reach the endpoint in ap-northeast-1.

✓ **C** Set the Regional API endpoint as an origin for the CloudFront distribution. Access the backend APIs by using the application's DNS name.
CloudFront can use a custom origin, and a cache behavior can route API requests to that origin. Having the application call its CloudFront-backed DNS name sends requests through CloudFront's edge network before they reach the Regional API endpoint, improving access for users outside ap-northeast-1. [1][2][3] [12]

https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_Origin.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/RequestAndResponseBehaviorCustomOrigin.html
https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_DefaultCacheBehavior.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistValuesOrigin.html

✗ **D** Set up AWS Global Accelerator in the same account that hosts the CloudFront distribution. Access the backend APIs by using Global Accelerator.
This does not configure CloudFront to route the API calls through its distribution, and Global Accelerator is not a direct way to accelerate a Regional API Gateway endpoint in this scenario. The application would not gain the proposed CloudFront path to the API.

Free demo · full exam: 681 questions

Question 52

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Recommend improvements to an existing solution to reduce costs.

A company recently deployed an application on Amazon ECS with the Fargate launch type. The company uses an AWS CloudFormation template to deploy the application. The company wants to identify over-provisioned resources and to update identified resources to optimize costs.

Which solution will meet these requirements?

A. Enable AWS Compute Optimizer. Identify resources that are classified as over-provisioned. Update the CloudFormation task definition to use the task size that Compute Optimizer recommends. Redeploy the CloudFormation template.

B. Enable AWS Compute Optimizer. Identify resources that are classified as over-provisioned. Implement an Amazon ECS target tracking scaling policy to dynamically adjust ECS services based on the task size that Compute Optimizer recommends.

C. Enable AWS Cost Explorer and use a Reserved Instance utilization and coverage report to identify over-provisioned ECS services. Create an Amazon ECS target tracking scaling policy to automatically adjust the services to match the values in the report.

D. Enable AWS Cost Explorer and use a Reserved Instance utilization and coverage report to identify over-provisioned ECS services. Update the CloudFormation task definition to use values that match the values in the report. Redeploy the CloudFormation template.

Answer: A. Enable AWS Compute Optimizer. Identify resources that are classified as over-provisioned. Update the CloudFormation task definition to use the task size that Compute Optimizer recommends. Redeploy the CloudFormation template.

Explanation

AWS Compute Optimizer provides ECS recommendations, including identifying over-provisioned Fargate task configurations. Updating the CloudFormation task definition to use the recommended task size and redeploying applies the optimization while keeping the application managed through CloudFormation. This directly addresses both identifying excess capacity and reducing costs by right-sizing the configured tasks.

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-recommendations.html
https://docs.aws.amazon.com/compute-optimizer/latest/ug/ecs-view-recommendations.html
https://docs.aws.amazon.com/compute-optimizer/latest/ug/view-ecs-recommendations.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/optimize-costs-microsoft-workloads/optimizer-ecs-fargate.html
https://docs.aws.amazon.com/compute-optimizer/latest/ug/supported-resources.html

✓ **A** Enable AWS Compute Optimizer. Identify resources that are classified as over-provisioned. Update the CloudFormation task definition to use the task size that Compute Optimizer recommends. Redeploy the CloudFormation template.

Compute Optimizer can provide ECS recommendations, including recommendations for Fargate task sizing and identifying over-provisioned resources. Changing the task definition to the recommended size and redeploying the CloudFormation template applies that right-sizing to the deployed application.

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-recommendations.html
https://docs.aws.amazon.com/compute-optimizer/latest/ug/ecs-view-recommendations.html
https://docs.aws.amazon.com/compute-optimizer/latest/ug/view-ecs-recommendations.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/optimize-costs-microsoft-workloads/optimizer-ecs-fargate.html
https://docs.aws.amazon.com/compute-optimizer/latest/ug/supported-resources.html

Free demo · full exam: 681 questions

✗ **B** Enable AWS Compute Optimizer. Identify resources that are classified as over-provisioned. Implement an Amazon ECS target tracking scaling policy to dynamically adjust ECS services based on the task size that Compute Optimizer recommends.
Compute Optimizer recommendations can inform changes to task sizing, but an ECS target tracking scaling policy adjusts the number of running tasks based on a metric; it does not change the task size in the task definition. This therefore does not apply the recommended right-sizing of the over-provisioned tasks.

https://docs.aws.amazon.com/compute-optimizer/latest/ug/ecs-view-recommendations.html
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/target-tracking-create-policy.html

✗ **C** Enable AWS Cost Explorer and use a Reserved Instance utilization and coverage report to identify over-provisioned ECS services. Create an Amazon ECS target tracking scaling policy to automatically adjust the services to match the values in the report.
Reserved Instance utilization and coverage reports concern reservation usage and do not provide the ECS task-sizing recommendations needed to identify over-provisioned Fargate tasks. Target tracking scales the number of tasks based on metrics, rather than right-sizing the task configuration.

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/target-tracking-create-policy.html

✗ **D** Enable AWS Cost Explorer and use a Reserved Instance utilization and coverage report to identify over-provisioned ECS services. Update the CloudFormation task definition to use values that match the values in the report. Redeploy the CloudFormation template.
A Reserved Instance utilization and coverage report does not identify recommended CPU and memory values for ECS task definitions. Although updating and redeploying the CloudFormation template is an appropriate way to apply task-definition changes, the report does not supply the right-sizing recommendations this scenario requires.

Question 53

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a strategy to improve overall operational excellence.

A company stores all its infrastructure code in a central Git repository that is hosted on GitHub Enterprise. The company wants to implement a process that submits every infrastructure update by using a pull request. Each pull request must automatically receive an AWS CloudFormation change set. After the pull request is approved, CloudFormation must apply the updates directly. The process must not require custom pipelines or manual template processing.

Which solution will meet these requirements?

A. Use CloudFormation to invoke an AWS CodePipeline action. Configure the CodePipeline action to validate the template and publish the validation results as pull request comments.

B. Use CloudFormation Git sync to link CloudFormation with the repository. Configure automatic change set creation for pull request updates. Configure CloudFormation to post the change set summary and apply the update when the pull request is approved and merged.

C. Use CloudFormation drift detection. Create an Amazon EventBridge rule that reacts to pull request updates. Configure the EventBridge rule to produce drift comparison results and send the results as comments before the stack is updated.

D. Use CloudFormation Lambda Hooks. Configure Lambda Hooks to evaluate the updated template and write the planned changes into the pull request when a stack update is requested.

Answer: B. Use CloudFormation Git sync to link CloudFormation with the repository. Configure automatic change set creation for pull request updates. Configure CloudFormation to post the change set summary and apply the update when the pull request is approved and merged.

Explanation

CloudFormation Git sync connects a CloudFormation stack to source code in a Git repository and supports pull-request workflows that generate change sets for review. That lets each infrastructure update be submitted through a pull request and gives reviewers a preview of the proposed stack changes. After approval and merge, CloudFormation can apply the update without requiring a custom pipeline or manual template processing.

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/git-sync-prereq.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/git-sync-create-stack-from-repository-source-code.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/git-sync.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **A** Use CloudFormation to invoke an AWS CodePipeline action. Configure the CodePipeline action to validate the template and publish the validation results as pull request comments.
This describes a separate CodePipeline-based validation workflow, rather than CloudFormation's native pull-request change-set workflow. It does not establish the required automatic change set and application of the approved update through CloudFormation Git sync.

✓ **B** Use CloudFormation Git sync to link CloudFormation with the repository. Configure automatic change set creation for pull request updates. Configure CloudFormation to post the change set summary and apply the update when the pull request is approved and merged.
CloudFormation Git sync links a stack to a repository and supports pull-request change-set review and deployment after a change is merged. This meets the requirements for automatic previews and applying approved infrastructure updates without custom pipelines or manual template processing.

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/git-sync-prereq.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/git-sync-create-stack-from-repository-source-code.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/git-sync.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **C** Use CloudFormation drift detection. Create an Amazon EventBridge rule that reacts to pull request updates. Configure the EventBridge rule to produce drift comparison results and send the results as comments before the stack is updated.
Drift detection compares a deployed stack with its actual resources to identify differences; it does not create a change set representing proposed pull-request updates. Adding EventBridge to send comparison results would not provide the requested native pull-request deployment workflow.

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/drift-aware-change-sets.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-stack-drift.html

✗ **D** Use CloudFormation Lambda Hooks. Configure Lambda Hooks to evaluate the updated template and write the planned changes into the pull request when a stack update is requested.
Lambda Hooks can run validation or policy logic during CloudFormation operations, but they are not the native mechanism for generating and presenting a pull-request change set. This approach would require custom hook behavior and does not meet the requirement to avoid custom processing.

https://docs.aws.amazon.com/cloudformation-cli/latest/hooks-userguide/lambda-hooks-create-lambda-function.html
https://docs.aws.amazon.com/cloudformation-cli/latest/hooks-userguide/creating-and-managing-hooks.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/TemplateReference/aws-resource-cloudformation-lambdahook.html

Free demo · full exam: 681 questions

Question 54

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a solution to improve security.

A company needs to improve the security of its web application on AWS. The application runs on a fleet of Amazon EC2 instances behind a public Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB is registered as a custom origin in an Amazon CloudFront distribution.

The company wants customers to access the website by using a fully qualified domain name (FQDN) that is associated with the CloudFront distribution. A security audit shows that the ALB can be accessed directly and that some requests bypass the CloudFront distribution.

The company needs a solution that will prevent direct access to the ALB. The solution also must ensure that all requests pass through the CloudFront distribution.

Which solution will meet these requirements?

A. Modify the security group that is attached to the ALB to allow incoming traffic from only the CloudFront distribution's origin.

B. Configure the CloudFront distribution to use an AWS WAF web ACL to block requests that bypass the CloudFront distribution.

C. Add a security header to the CloudFront distribution. Implement header validation and enforcement at the ALB.

D. Configure an origin access control (OAC) for the CloudFront distribution. Update the ALB to allow only requests that come from the OAC.

Answer: C. Add a security header to the CloudFront distribution. Implement header validation and enforcement at the ALB.

Explanation

Configure CloudFront to add a secret custom header to origin requests, then have the ALB validate that header and reject requests that do not contain it. This lets the ALB distinguish requests forwarded by the configured CloudFront distribution from direct requests, so customers can use the distribution's FQDN while direct ALB access is blocked. [8][11]

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/add-origin-custom-headers.html

✗ **A** Modify the security group that is attached to the ALB to allow incoming traffic from only the CloudFront distribution's origin.
An ALB security group cannot identify a particular CloudFront distribution by its origin; allowing CloudFront origin-facing IP ranges can permit requests from other distributions as well. It does not provide the distribution-specific validation needed to ensure requests pass through the intended distribution. [8]

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html

✗ **B** Configure the CloudFront distribution to use an AWS WAF web ACL to block requests that bypass the CloudFront distribution.
A web ACL associated with the CloudFront distribution evaluates requests that reach that distribution. It does not inspect requests sent directly to the public ALB, so it cannot block those bypass requests. [3]

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/IntroductionUseCases.html

✓ **C** Add a security header to the CloudFront distribution. Implement header validation and enforcement at the ALB.
CloudFront can add a custom header to requests it forwards to a custom origin, and the ALB can validate

Free demo · full exam: 681 questions

that header and reject requests without the expected value. Direct requests to the ALB will not have the required header, so this approach restricts access to requests forwarded by the configured distribution. [8][11]

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/add-origin-custom-headers.html

✗ **D** Configure an origin access control (OAC) for the CloudFront distribution. Update the ALB to allow only requests that come from the OAC.
OAC is not a mechanism that an ALB can use to authenticate CloudFront origin requests; it is intended for supported origins such as Amazon S3. For this ALB custom origin, a custom header validated at the ALB is the applicable approach. [6]

https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_Origin.html

Question 55

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Improve the reliability of existing solutions.

A company has implemented an ordering system using an event-driven architecture. During initial testing, the system stopped processing orders. Further log analysis revealed that one order message in an Amazon Simple Queue Service (Amazon SQS) standard queue was causing an error on the backend and blocking all subsequentorder messages The visibility timeout of the queue is set to 30 seconds, and the backend processing timeout is set to 10 seconds. A solutions architect needs to analyze faulty order messages and ensure that the system continues to process subsequent messages.

Which step should the solutions architect take to meet these requirements?

A. Increase the backend processing timeout to 30 seconds to match the visibility timeout.
B. Reduce the visibility timeout of the queue to automatically remove the faulty message.
C. Configure a new SQS FIFO queue as a dead-letter queue to isolate the faulty messages.
D. Configure a new SQS standard queue as a dead-letter queue to isolate the faulty messages.

Answer: D. Configure a new SQS standard queue as a dead-letter queue to isolate the faulty messages.

Explanation

Configure a standard SQS dead-letter queue and a redrive policy on the source standard queue. After a message repeatedly fails processing and exceeds the configured receive limit, SQS moves it to the dead-letter queue, where it can be analyzed separately while the source queue continues processing other orders. A standard source queue should use a standard dead-letter queue, not a FIFO queue.

https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-dead-letter-queue-redrive.html
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-dead-letter-queue.html?trk=article-ssr-frontend-pulse_little-text-block
https://docs.aws.amazon.com/en_gb/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-queue-types.html

✗ **A** Increase the backend processing timeout to 30 seconds to match the visibility timeout.
The backend timeout is already shorter than the 30-second visibility timeout, so an order that takes up to 10 seconds can finish before the message becomes visible again. Increasing the timeout does not isolate

Free demo · full exam: 681 questions

a message that causes a processing error; that faulty message could still be retried and interfere with processing.

https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/working-with-visibility-timeouts.html
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/best-practices-processing-messages-timely-manner.html

✗ **B** Reduce the visibility timeout of the queue to automatically remove the faulty message.
A visibility timeout does not automatically remove a message. When the timeout expires, an unacknowledged message becomes visible again and can be received and processed repeatedly, so reducing the timeout would not isolate the faulty order.

https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/working-with-visibility-timeouts.html
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-queue-parameters.html

✗ **C** Configure a new SQS FIFO queue as a dead-letter queue to isolate the faulty messages.
A dead-letter queue can isolate messages that repeatedly fail, but this source is an SQS standard queue and its dead-letter queue must also be standard. A FIFO dead-letter queue is therefore not the appropriate queue type for this source queue.

https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-dead-letter-queue-redrive.html
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-dead-letter-queue.html?trk=article-ssr-frontend-pulse_little-text-block
https://docs.aws.amazon.com/en_gb/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-queue-types.html

✓ **D** Configure a new SQS standard queue as a dead-letter queue to isolate the faulty messages.
Configure the standard queue as the source queue's dead-letter queue and set a redrive policy with an appropriate maximum receive count. Messages that repeatedly fail can then be moved out of the source queue for analysis, allowing subsequent order messages to be processed.

https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-dead-letter-queue-redrive.html
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-dead-letter-queue.html?trk=article-ssr-frontend-pulse_little-text-block
https://docs.aws.amazon.com/en_gb/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-queue-types.html

Dumpsgate

Free demo · full exam: 681 questions

Question 56

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a strategy to improve overall operational excellence.

A company is running a large containerized workload in the AWS Cloud. The workload consists of approximately 100 different services. The company uses Amazon ECS to orchestrate the workload.

Recently, the company's development team started using AWS Fargate instead of Amazon EC2 instances in the ECS cluster. In the past, the workload has come close to running the maximum number of EC2 instances that are available in the account.

The company is worried that the workload could reach the maximum number of ECS tasks that are allowed. A solutions architect must implement a solution that will notify the development team when Fargate reaches 80% of the maximum number of tasks.

What should the solutions architect do to meet this requirement?

A. Use Amazon CloudWatch to monitor the Sample Count statistic for each service in the ECS cluster. Set an alarm for when the math expression sample count/SERVICE_QUOTA service times 100 is greater than 80. Notify the development team by using Amazon SNS.

B. Use Amazon CloudWatch to monitor service quotas that are published under the AWS/Usage metric namespace. Set an alarm for when the math expression metric/SERVICE_QUOTA metric times 100 is greater than 80. Notify the development team by using Amazon SNS.

C. Create an AWS Lambda function to poll detailed metrics from the ECS cluster. When the number of running Fargate tasks is greater than 80, invoke Amazon SES to notify the development team.

D. Create an AWS Config rule to evaluate whether the Fargate service quota is greater than 80. Use Amazon SES to notify the development team when the AWS Config rule is not compliant.

Answer: B. Use Amazon CloudWatch to monitor service quotas that are published under the AWS/Usage metric namespace. Set an alarm for when the math expression metric/SERVICE_QUOTA metric times 100 is greater than 80. Notify the development team by using Amazon SNS.

Explanation

CloudWatch publishes AWS service usage metrics in the `AWS/Usage` namespace, which can be compared with the applicable service quota. A metric-math expression that divides usage by the quota and multiplies by 100 gives the percentage of the quota being used; an alarm at greater than 80% can notify the development team through Amazon SNS. This monitors quota utilization rather than relying on the EC2 instance limit or an arbitrary task-count threshold.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Usage-Metrics.html
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/using-metric-math.html
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Quotas-Visualize-Alarms.html
https://docs.aws.amazon.com/servicequotas/latest/userguide/configure-cloudwatch.html

✗ **A** Use Amazon CloudWatch to monitor the Sample Count statistic for each service in the ECS cluster. Set an alarm for when the math expression sample count/SERVICE_QUOTA service times 100 is greater than 80. Notify the development team by using Amazon SNS.
This uses a Sample Count statistic for ECS service metrics, rather than the CloudWatch usage metric that represents usage against the applicable service quota. Monitoring each service separately this way does not provide the required account-level Fargate quota utilization percentage. CloudWatch metric math and SNS alarms are useful components, but the metric selection is wrong for this requirement.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/using-metric-math.html
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/viewing_cloudwatch_metrics.html
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/available-metrics.html

Free demo · full exam: 681 questions

✓ **B** Use Amazon CloudWatch to monitor service quotas that are published under the AWS/Usage metric namespace. Set an alarm for when the math expression metric/SERVICE_QUOTA metric times 100 is greater than 80. Notify the development team by using Amazon SNS.
The `AWS/Usage` namespace provides usage metrics for AWS services that can be compared with their quotas. Dividing the usage metric by the `SERVICE_QUOTA` value and multiplying by 100 calculates the percentage of the quota in use, so a CloudWatch alarm can notify the team when usage exceeds 80%.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Usage-Metrics.html
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/using-metric-math.html
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Quotas-Visualize-Alarms.html
https://docs.aws.amazon.com/servicequotas/latest/userguide/configure-cloudwatch.html

✗ **C** Create an AWS Lambda function to poll detailed metrics from the ECS cluster. When the number of running Fargate tasks is greater than 80, invoke Amazon SES to notify the development team.
A threshold of 80 running tasks is an absolute count, not 80% of the maximum allowed tasks; the quota may be higher or lower. Polling ECS metrics with a custom Lambda function also does not directly compare Fargate usage with its service quota as the requirement calls for.

https://docs.aws.amazon.com/general/latest/gr/ecs-service.html
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/service_utilization.html
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/viewing_cloudwatch_metrics.html
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/available-metrics.html

✗ **D** Create an AWS Config rule to evaluate whether the Fargate service quota is greater than 80. Use Amazon SES to notify the development team when the AWS Config rule is not compliant.
The requirement is to notify the team when Fargate usage reaches 80% of its quota, not when the quota value itself is greater than 80. An AWS Config rule is not the CloudWatch usage-metric alarm needed to calculate and alert on that utilization percentage.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Usage-Metrics.html
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Quotas-Visualize-Alarms.html
https://docs.aws.amazon.com/servicequotas/latest/userguide/configure-cloudwatch.html

---

Question 57

Multiple choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a solution to improve security.

A company is storing sensitive data in an Amazon S3 bucket. The company must log all activities for objects in the S3 bucket and must keep the logs for 5 years. The company's security team also must receive an email notification every time there is an attempt to delete data in the S3 bucket.

Which combination of steps will meet these requirements MOST cost-effectively? (Select THREE.)

A. Configure AWS CloudTrail to log S3 data events.

B. Configure S3 server access logging for the S3 bucket.

C. Configure Amazon S3 to send object deletion events to Amazon Simple Email Service (Amazon SES).

D. Configure Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic.

E. Configure Amazon S3 to send the logs to Amazon Timestream with data storage tiering.

F. Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.

Free demo · full exam: 681 questions

Answer

- A. **Configure AWS CloudTrail to log S3 data events.**
- D. **Configure Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic.**
- F. **Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.**

Explanation

CloudTrail data events provide object-level activity logging for the S3 bucket. EventBridge can route object deletion events to an SNS topic, which can email the security team. Storing the logs in a separate S3 bucket with a lifecycle policy supports the five-year retention requirement while allowing storage costs to be managed through lifecycle transitions.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/enable-cloudtrail-logging-for-s3.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/cloudtrail-logging-s3-info.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-log-s3-data-events.html
https://docs.aws.amazon.com/eventbridge/latest/ref/events-ref-s3.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/deleting-log-files-lifecycle.html

✓ **A** Configure AWS CloudTrail to log S3 data events.
CloudTrail data events record object-level S3 activity, which meets the requirement to log activities involving the bucket's objects. The resulting logs can be delivered to the separate S3 bucket for retention.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/enable-cloudtrail-logging-for-s3.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/cloudtrail-logging-s3-info.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/cloudtrail-logging.html

✗ **B** Configure S3 server access logging for the S3 bucket.
Server access logging produces request logs, but it is not the selected mechanism for capturing the required object-level activity in this solution. CloudTrail data events are the appropriate choice for that requirement.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerLogs.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/enable-cloudtrail-logging-for-s3.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/cloudtrail-logging-s3-info.html

✗ **C** Configure Amazon S3 to send object deletion events to Amazon Simple Email Service (Amazon SES).
Amazon S3 event notifications do not send directly to SES as a destination. Use a supported event-routing service such as EventBridge and an SNS topic to deliver email notifications instead.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/EventNotifications.html

✓ **D** Configure Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic.
EventBridge can receive S3 object deletion events and route them to an SNS topic. The security team can subscribe by email to that topic to receive notifications when deletion events occur.

https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-log-s3-data-events.html
https://docs.aws.amazon.com/eventbridge/latest/ref/events-ref-s3.html

✗ **E** Configure Amazon S3 to send the logs to Amazon Timestream with data storage tiering.
Timestream is not the log-storage destination used in the selected design. A separate S3 bucket with a lifecycle policy is the cost-conscious way to retain these logs for the required period.

✓ **F** Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.
A separate bucket provides a destination for the logs, and its lifecycle policy can manage storage over the five-year retention period. Lifecycle transitions can help reduce the cost of keeping logs as they age.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/troubleshoot-lifecycle.html
https://docs.aws.amazon.com/AmazonS3/latest/userguide/deleting-log-files-lifecycle.html

Free demo · full exam: 681 questions

Question 58

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine how to improve the performance of existing solutions.

A company hosts an application on AWS. The application uses AWS Lambda functions that are invoked by an Amazon API Gateway API. The company has an Amazon CloudFront distribution that uses the API Gateway API as its origin. The CloudFront distribution serves web requests to customers worldwide.

During testing, users experienced slow responses from the application APIs. The company discovered that requests from different AWS Regions contained inconsistent query parameters with mixed-case letters, which caused increased cache misses and more requests to reach the Lambda functions.

The company wants to ensure that the API consistently provides responses with minimal latency.

Which solution will meet these requirements?

A. Create a new Lambda function to sort incoming request query parameters alphabetically and convert the parameters to lowercase. Configure the CloudFront distribution to use the Lambda@Edge function type. Configure the Lambda function to invoke on origin request.

B. Create a CloudFront function to sort incoming request query parameters alphabetically and convert the parameters to lowercase. Configure the CloudFront distribution to use the CloudFront Functions function type. Configure the CloudFront function to invoke on viewer request.

C. Configure the API Gateway API to use mapping templates to sort incoming request query parameters alphabetically and convert the parameters to lowercase before Lambda processes the request.

D. Configure the API Gateway API to use a Lambda authorizer to sort incoming request query parameters alphabetically and convert the parameters to lowercase before Lambda processes the request.

Answer: B. Create a CloudFront function to sort incoming request query parameters alphabetically and convert the parameters to lowercase. Configure the CloudFront distribution to use the CloudFront Functions function type. Configure the CloudFront function to invoke on viewer request.

Explanation

A CloudFront function on the viewer-request event can normalize query parameter names and ordering before CloudFront performs its cache lookup. That makes equivalent requests use a consistent cache key, reducing cache misses and requests that reach API Gateway and Lambda, which supports the goal of minimizing latency for customers worldwide. An origin-request function runs only after a cache miss, so it would not prevent the misses.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/example_cloudfront_functions_normalize_query_string_parameters_section.html
https://docs.aws.amazon.com/pdfs/AmazonCloudFront/latest/DeveloperGuide/AmazonCloudFront_DevGuide.pdf
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/functions-event-structure.html

✗ **A** Create a new Lambda function to sort incoming request query parameters alphabetically and convert the parameters to lowercase. Configure the CloudFront distribution to use the Lambda@Edge function type. Configure the Lambda function to invoke on origin request.

Free demo · full exam: 681 questions

An origin-request function runs after CloudFront has determined that the request is a cache miss. It therefore cannot normalize the request before cache-key evaluation to prevent the misses described in the scenario. Lambda@Edge would also be unnecessary when a CloudFront function can perform this viewer-request normalization.

https://docs.aws.amazon.com/pdfs/AmazonCloudFront/latest/DeveloperGuide/AmazonCloudFront_DevGuide.pdf
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-event-structure.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html

✓ **B** Create a CloudFront function to sort incoming request query parameters alphabetically and convert the parameters to lowercase. Configure the CloudFront distribution to use the CloudFront Functions function type. Configure the CloudFront function to invoke on viewer request.
The viewer-request event occurs before CloudFront checks its cache, so the function can normalize the query string before it affects the cache key. Sorting parameters and converting them to lowercase makes equivalent requests consistent, helping reduce cache misses and avoid unnecessary requests to API Gateway and Lambda.

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/example_cloudfront_functions_normalize_query_string_parameters_section.html
https://docs.aws.amazon.com/pdfs/AmazonCloudFront/latest/DeveloperGuide/AmazonCloudFront_DevGuide.pdf
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/functions-event-structure.html

✗ **C** Configure the API Gateway API to use mapping templates to sort incoming request query parameters alphabetically and convert the parameters to lowercase before Lambda processes the request.
API Gateway mapping templates operate after the request has passed through CloudFront. They cannot normalize the query string before CloudFront's cache lookup, so the inconsistent cache keys and resulting misses would remain. The approach therefore does not address the source of the increased latency in this scenario.

https://docs.aws.amazon.com/apigateway/latest/developerguide/request-response-data-mappings.html
https://docs.aws.amazon.com/apigateway/latest/developerguide/rest-api-parameter-mapping-sources.html

✗ **D** Configure the API Gateway API to use a Lambda authorizer to sort incoming request query parameters alphabetically and convert the parameters to lowercase before Lambda processes the request.
A Lambda authorizer is part of API Gateway request processing, which occurs after a request has already reached the API origin. It cannot normalize parameters before CloudFront evaluates the cache key, and adding authorizer processing does not prevent the described cache misses.

https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-lambda-authorizer.html

Question 59

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a solution to improve security.

A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.

Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts Which solution meets these requirements with the LEAST amount of operational overhead?

A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organization. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage Update the parameter as needed to add or remove accounts or OUs Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account

B. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rules. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resources Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.

C. Create AWS WAF rules in the management account of the organization Use AWS Lambda environment variables to store account numbers and OUs to manage Update environment variables as needed to add or remove accounts or OUs Create cross-account IAM roles in member accounts Assume the rotes by using AWS Security Token Service (AWS STS) in the Lambda function tocreate and update AWS WAF rules in the member accounts.

D. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage Update AWS KMS as needed to add or remove accounts or OUs Create IAM users in member accounts Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organization. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage Update the parameter as needed to add or remove accounts or OUs Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account

Explanation

AWS Firewall Manager is designed to centrally manage AWS WAF policies across accounts in an AWS Organization, including applying policies to selected accounts or OUs and remediating noncompliant resources. Storing the target accounts and OUs in Parameter Store gives administrators a place to adjust the managed scope, while the EventBridge and Lambda workflow updates the Firewall Manager policy when that list changes. This meets the requested flexibility and automated remediation without building a separate cross-account WAF management system.

https://docs.aws.amazon.com/waf/latest/developerguide/working-with-policies.html
https://docs.aws.amazon.com/waf/latest/developerguide/join-aws-orgs.html

✓ **A** Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organization. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage.

Free demo · full exam: 681 questions

Update the parameter as needed to add or remove accounts or OUs. Use an Amazon EventBridge rule to identify changes to the parameter and invoke a Lambda function to update the security policy in the Firewall Manager administrative account.
Firewall Manager provides centralized AWS WAF policy management across an organization, and its policies can be scoped to accounts and OUs and used to remediate noncompliant resources. The parameter and event-driven function let administrators change the target scope and have the central policy updated, meeting the requirements with less ongoing work than managing WAF rules separately in each account.

https://docs.aws.amazon.com/waf/latest/developerguide/working-with-policies.html
https://docs.aws.amazon.com/waf/latest/developerguide/join-aws-orgs.html

✗ **B** Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rules. Deploy automated remediation actions by using Lambda to fix noncompliant resources. Deploy AWS WAF rules by using a CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
This approach combines Config rules, Lambda remediation, and StackSets rather than using Firewall Manager's centralized WAF policy management and remediation capabilities. It adds deployment and maintenance components, so it does not meet the least-operational-overhead requirement as well as option A.

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-orgs-associate-stackset-with-org.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/TemplateReference/aws-properties-cloudformation-stackset-deploymenttargets.html
https://docs.aws.amazon.com/waf/latest/developerguide/working-with-policies.html

✗ **C** Create AWS WAF rules in the management account of the organization. Use Lambda environment variables to store account numbers and OUs to manage. Update environment variables as needed to add or remove accounts or OUs. Create cross-account IAM roles in member accounts and assume the roles by using STS in the Lambda function to create and update AWS WAF rules in the member accounts.
This requires custom Lambda logic to assume roles and create or update WAF rules across member accounts, as well as cross-account roles that must be maintained. It lacks the centralized policy and remediation management offered by Firewall Manager and therefore creates more operational overhead.

https://docs.aws.amazon.com/waf/latest/developerguide/working-with-policies.html
https://docs.aws.amazon.com/waf/latest/developerguide/join-aws-orgs.html

✗ **D** Use AWS Control Tower to manage AWS WAF rules across accounts in the organization. Use AWS KMS to store account numbers and OUs to manage. Update AWS KMS as needed to add or remove accounts or OUs. Create IAM users in member accounts and allow AWS Control Tower in the management account to use access keys and secret access keys to create and update AWS WAF rules in the member accounts.
Control Tower is not the AWS service for centrally managing and remediating AWS WAF policies across accounts; Firewall Manager is intended for that purpose. Storing targeting data in KMS and using IAM users' long-term access keys for cross-account updates also adds unnecessary complexity and credential-management risk.

https://docs.aws.amazon.com/waf/latest/developerguide/working-with-policies.html
https://docs.aws.amazon.com/waf/latest/developerguide/join-aws-orgs.html

Free demo · full exam: 681 questions

Question 60

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Improve the reliability of existing solutions.

A company has a website that serves many visitors. The company deploys a backend service for the website in a primary AWS Region and a disaster recovery (DR) Region.

A single Amazon CloudFront distribution is deployed for the website. The company creates an Amazon Route 53 record set with health checks and a failover routing policy for the primary Region's backend service. The company configures the Route 53 record set as an origin for the CloudFront distribution. The company configures another record set that points to the backend service's endpoint in the DR Region as a secondary failover record type. The TTL for both record sets is 60 seconds.

Currently, failover takes more than 1 minute. A solutions architect must design a solution that will provide the fastest failover time.

Which solution will achieve this goal?

A. Deploy an additional CloudFront distribution. Create a new Route 53 failover record set with health checks for both CloudFront distributions.

B. Set the TTL to 1 second for the existing Route 53 record sets that are used for the backend service in each Region.

C. Create new record sets for the backend services by using a latency routing policy. Use the record sets as an origin in the CloudFront distribution.

D. Create a CloudFront origin group that includes two origins, one for each backend service Region. Configure origin failover as a cache behavior for the CloudFront distribution.

Answer: D. Create a CloudFront origin group that includes two origins, one for each backend service Region. Configure origin failover as a cache behavior for the CloudFront distribution.

Explanation

A CloudFront origin group can contain the primary and DR backend origins and fail over directly to the secondary when the primary returns a configured failure status code. This avoids relying on Route 53 health-check detection and DNS TTL expiration to redirect requests, addressing the requirement for the fastest failover through the existing CloudFront distribution. [1][4][5]

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/disaster-recovery-resiliency.html
https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_OriginGroup.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/RequestAndResponseBehaviorOriginGroups.html

✗ **A** Deploy an additional CloudFront distribution. Create a new Route 53 failover record set with health checks for both CloudFront distributions.
This adds another layer of DNS-based failover rather than making the existing distribution switch backend origins directly. Route 53 health checks and DNS caching can delay when clients use the alternate distribution, so this does not provide the fastest failover. [2][3][8]

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-values-failover.html
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-problems.html
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-failover.html

✗ **B** Set the TTL to 1 second for the existing Route 53 record sets that are used for the backend service in each Region.
A shorter TTL can reduce DNS caching delays, but it does not eliminate the time needed for health checks

Free demo · full exam: 681 questions

and DNS-based failover to take effect. CloudFront origin failover switches between configured origins without depending on clients receiving an updated DNS answer. [3][5]

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-problems.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/RequestAndResponseBehaviorOriginGroups.html

✗ **C** Create new record sets for the backend services by using a latency routing policy. Use the record sets as an origin in the CloudFront distribution.
Latency-based routing directs DNS queries to an endpoint based on latency; it is not a direct, request-time failover mechanism between CloudFront origins. It therefore does not address the delay inherent in the current DNS-based failover design. [2][8]

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-values-failover.html
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-failover.html

✓ **D** Create a CloudFront origin group that includes two origins, one for each backend service Region. Configure origin failover as a cache behavior for the CloudFront distribution.
CloudFront origin groups support a primary origin and a secondary origin, and CloudFront can retry a request against the secondary when the primary returns a configured failure status code. This lets the existing distribution fail over between regional backends directly, without waiting for Route 53 health checks or TTL-based DNS updates. [1][4][5]

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/disaster-recovery-resiliency.html
https://docs.aws.amazon.com/cloudfront/latest/APIReference/API_OriginGroup.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/RequestAndResponseBehaviorOriginGroups.html

Dumpsgate

Question 61

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Improve the reliability of existing solutions.

A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes.

The company must maintain backups in a separate AWS Region. The company must be able to recover the EC2 instances and their configuration within I business day, with loss of no more than I day's worth of data. The company has limited staff and needs a backup solution that optimizes operational efficiency and cost. The company already has created an AWS CloudFormation template that can deploy the required network configuration in a secondary Region.

Which solution will meet these requirements?

A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbooks. Copy the snapshots to the secondary Region. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.

B. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volumes. In the event of a failure, launch theCloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.

C. Use AWS Backup to create a scheduled daily backup plan for the EC2 instances. Configure the backup task to copy the backups to a vault in the secondary Region. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.

D. Deploy EC2 instances of the same size and configuration to the secondary Region. Configure AWS DataSync daily to copy data from the primary Region to the secondary Region. In the event of a failure, launch the CloudFormation template and transfer usage to the secondaryRegion.

Answer: C. Use AWS Backup to create a scheduled daily backup plan for the EC2 instances. Configure the backup task to copy the backups to a vault in the secondary Region. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.

Explanation

C is correct because AWS Backup can create scheduled daily backups of the EC2 instances and copy recovery points to a vault in another Region. That supports the one-day data-loss limit while keeping the backups available for a regional recovery. AWS Backup can restore the EC2 instances and their configuration, and the existing CloudFormation template can provide the secondary-Region network. This avoids building and maintaining a separate snapshot-and-recovery workflow, which fits the company's limited staff and operational-efficiency requirement.

https://docs.aws.amazon.com/aws-backup/latest/devguide/restoring-ec2.html
https://docs.aws.amazon.com/aws-backup/latest/devguide/recov-point-create-a-copy.html
https://docs.aws.amazon.com/aws-backup/latest/devguide/cross-region-backup.html

✗ **A** Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbooks. Copy the snapshots to the secondary Region. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.

Free demo · full exam: 681 questions

This approach can provide daily, cross-Region snapshots, but it requires a separate template and runbook-based snapshot and copy workflow in addition to restoring the volumes. That is more custom operational work than using AWS Backup to manage scheduled EC2 backups, copies, and instance recovery, which is a better fit for the limited staff.

https://docs.aws.amazon.com/systems-manager-automation-runbooks/latest/userguide/automation-aws-copysnapshot.html
https://docs.aws.amazon.com/ebs/latest/userguide/ebs-copy-snapshot.html

✗ **B** Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volumes. In the event of a failure, launch the CloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.
Amazon DLM can automate EBS snapshot lifecycle policies, but snapshots alone do not provide the same EC2 instance-and-configuration restore workflow as AWS Backup. The recovery would still require additional steps to recreate the instances and restore their volumes, making this a less complete and less operationally efficient solution for the stated recovery requirement.

https://docs.aws.amazon.com/aws-backup/latest/devguide/restoring-ec2.html
https://docs.aws.amazon.com/ebs/latest/userguide/dlm-elements.html

✓ **C** Use AWS Backup to create a scheduled daily backup plan for the EC2 instances. Configure the backup task to copy the backups to a vault in the secondary Region. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.
AWS Backup supports scheduled EC2 backups and copying recovery points to another Region, satisfying the daily recovery-point objective and regional-backup requirement. It can restore EC2 instances with their configuration, while the existing CloudFormation template supplies the secondary-Region network. This managed workflow reduces the custom backup and recovery operations the company would otherwise have to maintain.

https://docs.aws.amazon.com/aws-backup/latest/devguide/restoring-ec2.html
https://docs.aws.amazon.com/aws-backup/latest/devguide/recov-point-create-a-copy.html
https://docs.aws.amazon.com/aws-backup/latest/devguide/cross-region-backup.html

✗ **D** Deploy EC2 instances of the same size and configuration to the secondary Region. Configure AWS DataSync daily to copy data from the primary Region to the secondary Region. In the event of a failure, launch the CloudFormation template and transfer usage to the secondary Region.
Copying data daily does not by itself create recoverable EC2 backups that preserve the instances' attached EBS volumes and configuration. The proposal also requires pre-deployed secondary instances, which does not eliminate the need for a backup-and-restore method for the persistent data. It therefore does not meet the stated recovery requirements as completely as AWS Backup.

Free demo · full exam: 681 questions

Question 62

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a strategy to improve overall operational excellence.

A company is using an organization in AWS organization to manage AWS accounts. For each new project the company creates a new linked account. After the creation of a new account, the root user signs in to the new account and creates a service request to increase the service quota for Amazon EC2 instances. A solutions architect needs to automate this process.

Which solution will meet these requirements with tie LEAST operational overhead?

A. Create an Amazon EventBridge rule to detect creation of a new account Send the event to an Amazon Simple Notification Service (Amazon SNS) topic that invokes an AWS Lambda function. Configure the Lambda function to run the request-service-quota-increase command to request a service quota increase for EC2 instances.

B. Create a Service Quotas request template in the management account. Configure the desired service quota increases for EC2 instances.

C. Create an AWS Config rule in the management account to set the service quota for EC2 instances.

D. Create an Amazon EventBridge rule to detect creation of a new account. Send the event to an Amazon simple Notification service (Amazon SNS) topic that involves an AWS Lambda function. Configure the Lambda function to run the create-case command to request a service quota increase for EC2 instances.

Answer: B. Create a Service Quotas request template in the management account. Configure the desired service quota increases for EC2 instances.

Explanation

A Service Quotas request template in the organization's management account lets the company configure the desired EC2 quota increases centrally and have the requests applied for new accounts. This avoids building and maintaining an event-driven Lambda workflow for every account creation, so it has the least operational overhead. [1]

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html

✗ **A** Create an Amazon EventBridge rule to detect creation of a new account. Send the event to an Amazon SNS topic that invokes an AWS Lambda function. Configure the Lambda function to run the request-service-quota-increase command to request a service quota increase for EC2 instances.
The CLI command can request a quota increase, but this approach requires maintaining an EventBridge rule, SNS topic, Lambda function, and its permissions and logic. A Service Quotas request template handles the organization's new-account use case with less operational work. [1][2]

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html
https://docs.aws.amazon.com/cli/latest/reference/service-quotas/request-service-quota-increase.html

✓ **B** Create a Service Quotas request template in the management account. Configure the desired service quota increases for EC2 instances.
Service Quotas request templates let an organization configure quota increase requests centrally for new accounts. This avoids requiring the root user to sign in and submit a request in each account, and avoids custom automation. [1]

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html

✗ **C** Create an AWS Config rule in the management account to set the service quota for EC2 instances.
AWS Config rules evaluate resource configuration and can report or remediate certain configuration issues; they are not the mechanism for requesting service quota increases. Service Quotas request

Free demo · full exam: 681 questions

templates are designed for centrally configuring quota requests for organization accounts. [1][7]

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html
https://docs.aws.amazon.com/config/latest/developerguide/configlimits.html

✗ **D** Create an Amazon EventBridge rule to detect creation of a new account. Send the event to an Amazon SNS topic that invokes an AWS Lambda function. Configure the Lambda function to run the create-case command to request a service quota increase for EC2 instances.
This still requires the event-driven SNS and Lambda infrastructure, so it has more operational overhead than a request template. The option also uses a support-case workflow rather than the organization's centrally configured Service Quotas request-template approach. [1][3]

https://docs.aws.amazon.com/servicequotas/latest/userguide/organization-templates.html
https://docs.aws.amazon.com/awssupport/latest/user/create-service-quota-increase.html

Question 63

Multiple choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a solution to improve security.

A company has several AWS Lambda functions written in Python. The functions are deployed with the .zip package deployment type. The functions use a Lambda layer that contains common libraries and packages in a .zip file. The Lambda .zip packages and the Lambda layer .zip file are stored in an Amazon S3 bucket.

The company must implement automatic scanning of the Lambda functions and the Lambda layer to identify CVEs. A subset of the Lambda functions must receive automated code scans to detect potential data leaks and other vulnerabilities. The code scans must occur only for selected Lambda functions, not all the Lambda functions.

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

A. Activate Amazon Inspector. Start automated CVE scans.

B. Activate Lambda standard scanning and Lambda code scanning in Amazon Inspector.

C. Enable Amazon GuardDuty. Enable the Lambda Protection feature in GuardDuty.

D. Enable scanning in the Monitor settings of the Lambda functions that need code scans.

E. Tag Lambda functions that do not need code scans. In the tag, include a key of InspectorCodeExclusion and a value of LambdaCodeScanning.

F. Use Amazon Inspector to scan the S3 bucket that contains the Lambda .zip packages and the Lambda layer .zip file for code scans.

Answer

- B. **Activate Lambda standard scanning and Lambda code scanning in Amazon Inspector.**
- D. **Enable scanning in the Monitor settings of the Lambda functions that need code scans.**
- E. **Tag Lambda functions that do not need code scans. In the tag, include a key of InspectorCodeExclusion and a value of LambdaCodeScanning.**

Explanation

Amazon Inspector's Lambda standard scanning identifies CVEs in function dependencies, while Lambda

Free demo · full exam: 681 questions

code scanning detects issues such as potential data leaks in function code. Enable both scan types, then enable code scanning for the functions that need it in their Monitor settings. Tag functions that should not receive code scans with the specified exclusion tag; this keeps code scanning limited to the selected functions. [1, 2, 3, 6]

https://docs.aws.amazon.com/inspector/latest/user/scanning-lambda.html
https://docs.aws.amazon.com/inspector/latest/user/scanning_resources_lambda.html
https://docs.aws.amazon.com/inspector/latest/user/scanning_resources_lambda_code_exclude_functions.html
https://docs.aws.amazon.com/inspector/latest/user/scanning_resources_lambda_code.html

✗ **A** Activate Amazon Inspector. Start automated CVE scans.
Inspector standard scanning can identify CVEs, but this action alone does not enable the code scans needed to detect potential data leaks and other code vulnerabilities. The required combination also configures both Lambda scan types and limits code scans to selected functions.

https://docs.aws.amazon.com/inspector/latest/user/scanning-lambda.html
https://docs.aws.amazon.com/inspector/latest/user/scanning_resources_lambda.html
https://docs.aws.amazon.com/inspector/latest/user/activate-scans.html

✓ **B** Activate Lambda standard scanning and Lambda code scanning in Amazon Inspector.
Lambda standard scanning checks for CVEs in function dependencies, and Lambda code scanning analyzes function code for security issues such as potential data leaks. Activating both scan types addresses both scanning requirements; the other actions control which functions receive code scans.

https://docs.aws.amazon.com/inspector/latest/user/scanning-lambda.html
https://docs.aws.amazon.com/inspector/latest/user/scanning_resources_lambda.html
https://docs.aws.amazon.com/inspector/latest/user/scanning_resources_lambda_code.html

✗ **C** Enable Amazon GuardDuty. Enable the Lambda Protection feature in GuardDuty.
GuardDuty Lambda Protection monitors Lambda network activity to identify threats; it is not the Inspector scanning configuration for CVEs in Lambda dependencies or code-level vulnerabilities such as data leaks. It also does not provide the required selection of functions for code scans.

https://docs.aws.amazon.com/guardduty/latest/ug/lambda-protection.html
https://docs.aws.amazon.com/guardduty/latest/ug/remediate-lambda-protection-finding-types.html

✓ **D** Enable scanning in the Monitor settings of the Lambda functions that need code scans.
Enabling code scanning in the Monitor settings of the chosen functions applies code scans to those functions. This meets the requirement to scan only a subset rather than every Lambda function.

https://docs.aws.amazon.com/inspector/latest/user/scanning-lambda.html
https://docs.aws.amazon.com/inspector/latest/user/scanning_resources_lambda_code.html

✓ **E** Tag Lambda functions that do not need code scans. In the tag, include a key of InspectorCodeExclusion and a value of LambdaCodeScanning.
The InspectorCodeExclusion:LambdaCodeScanning tag excludes a function from Lambda code scanning. Applying it to functions that do not need code scans leaves the selected functions eligible for code scanning while excluding the others.

https://docs.aws.amazon.com/inspector/latest/user/scanning_resources_lambda_code_exclude_functions.html

✗ **F** Use Amazon Inspector to scan the S3 bucket that contains the Lambda .zip packages and the Lambda layer .zip file for code scans.
The requirement is to scan the deployed Lambda functions and layer, not to run Inspector code scans against the S3 bucket containing their deployment archives. Inspector's Lambda scanning applies to supported Lambda resources, so this does not configure the required function-level code scans or their selection.

https://docs.aws.amazon.com/inspector/latest/user/scanning-lambda.html
https://docs.aws.amazon.com/inspector/latest/user/scanning_resources_lambda.html
https://docs.aws.amazon.com/inspector/latest/user/scanning-resources.html
https://docs.aws.amazon.com/inspector/latest/user/supported.html

Free demo · full exam: 681 questions

Question 64

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Improve the reliability of existing solutions.

A company has a website that runs on four Amazon EC2 instances that are behind an Application Load Balancer (ALB). When the ALB detects that an EC2 instance is no longer available, an Amazon CloudWatch alarm enters the ALARM state. A member of the company's operations team then manually adds a new EC2 instance behind the ALB.

A solutions architect needs to design a highly available solution that automatically handles the replacement of EC2 instances. The company needs to minimize downtime during the switch to the new solution.

Which set of steps should the solutions architect take to meet these requirements?

A. Delete the existing ALB. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Create a new ALB. Attach the Auto Scaling group to the new ALB. Attach the existing EC2 instances to the Auto Scaling group.

B. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Attach the Auto Scaling group to the existing ALB. Attach the existing EC2 instances to the Auto Scaling group.

C. Delete the existing ALB and the EC2 instances. Create an Auto Scaling group that is configuredto handle the web application traffic. Attach a new launch template to the Auto Scaling group. Create a new ALB. Attach the Auto Scaling group to the new ALB. Wait for the Auto Scaling group to launch the minimum number of EC2 instances.

D. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Attach the Auto Scaling group to the existing ALB. Wait for the existing ALB to register the existing EC2 instances with the Auto Scaling group.

Answer: B. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Attach the Auto Scaling group to the existing ALB. Attach the existing EC2 instances to the Auto Scaling group.

Explanation

B preserves the existing ALB and running EC2 instances, minimizing disruption while the solution is migrated. The architect can create an Auto Scaling group with a launch template, associate it with the existing ALB, and attach the current instances to the group. The Auto Scaling group can then manage capacity and replace unhealthy instances automatically, rather than relying on an operations-team member to add replacements manually.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/health-checks-overview.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-health-checks.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-detach-attach-instances.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/create-asg-launch-template.html

✗ **A** Delete the existing ALB. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Create a new ALB. Attach the Auto Scaling group to the new ALB. Attach the existing EC2 instances to the Auto Scaling group. Deleting and replacing the existing ALB creates unnecessary risk and potential downtime, contrary to the requirement to minimize disruption. The existing ALB can be associated with the Auto Scaling group, and the running instances can be attached to that group instead.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-detach-attach-instances.html

✓ **B** Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Attach the Auto Scaling group to the existing ALB. Attach the existing EC2 instances to the Auto Scaling group.
This keeps the existing ALB and EC2 instances in service while bringing them under Auto Scaling management. With the ALB associated with the group and the existing instances attached, Auto Scaling can use its health-check and replacement capabilities to automate recovery when an instance becomes unhealthy.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/health-checks-overview.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-health-checks.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-detach-attach-instances.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/create-asg-launch-template.html

✗ **C** Delete the existing ALB and the EC2 instances. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Create a new ALB. Attach the Auto Scaling group to the new ALB. Wait for the Auto Scaling group to launch the minimum number of EC2 instances.
Deleting all existing instances and the ALB before the replacement resources are ready would cause avoidable downtime. This approach does not preserve the running service during the transition, unlike using the existing ALB and attaching the current instances to the Auto Scaling group.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-detach-attach-instances.html

✗ **D** Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Attach the Auto Scaling group to the existing ALB. Wait for the existing ALB to register the existing EC2 instances with the Auto Scaling group.
The ALB does not register existing EC2 instances with the Auto Scaling group; the instances must be attached to the group. Without that attachment, the existing instances are not brought under Auto Scaling management as intended.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-detach-attach-instances.html

Dumpsgate

Question 65

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a solution to improve security.

A company is running an application that uses an Amazon ElastiCache for Redis cluster as a caching layer A recent security audit revealed that the company has configured encryption at rest for ElastiCache However the company did not configure ElastiCache to use encryption in transit Additionally, users can access the cache without authentication A solutions architect must make changes to require user authentication and to ensure that the company is using end-to-end encryption Which solution will meet these requirements?

A. Create an AUTH token Store the token in AWS System Manager Parameter Store, as anencrypted parameter Create a new cluster with AUTH and configure encryption in transit Update the application to retrieve the AUTH token from Parameter Store when necessary and to use the AUTH token for authentication

B. Create an AUTH token Store the token in AWS Secrets Manager Configure the existing cluster to use the AUTH token and configure encryption in transit Update the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication.

C. Create an SSL certificate Store the certificate in AWS Secrets Manager Create a new cluster and configure encryption in transit Update the application to retrieve the SSL certificate from Secrets Manager when necessary and to use the certificate for authentication.

D. Create an SSL certificate Store the certificate in AWS Systems Manager Parameter Store, as an encrypted advanced parameter Update the existing cluster to configure encryption in transit Update the application to retrieve the SSL certificate from Parameter Store when necessary and to use the certificate for authentication

Answer: B. Create an AUTH token Store the token in AWS Secrets Manager Configure the existing cluster to use the AUTH token and configure encryption in transit Update the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication.

Explanation

Option B meets both requirements by configuring the Redis cluster to require an AUTH token and enabling encryption in transit, so clients authenticate and connections are encrypted. Storing the token in AWS Secrets Manager lets the application retrieve the credential when it needs to connect. The other options either use an SSL certificate as though it were a Redis authentication credential or do not configure the existing cluster as described in the correct solution.

https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/encryption.html
https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/auth.html
https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/in-transit-encryption.html

✗ **A** Create an AUTH token, store it in AWS Systems Manager Parameter Store as an encrypted parameter, create a new cluster with AUTH and encryption in transit, and update the application to retrieve and use the token.
An AUTH token and encryption in transit address authentication and encrypted connections, but this option creates a new cluster rather than configuring the existing cluster as in the correct solution. It also uses Parameter Store rather than the specified Secrets Manager approach.

https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/encryption.html
https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/auth.html
https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/in-transit-encryption.html

✓ **B** Create an AUTH token, store it in AWS Secrets Manager, configure the existing cluster to use the

Free demo · full exam: 681 questions

token and encryption in transit, and update the application to retrieve and use the token.
Redis AUTH provides the required user authentication, while encryption in transit protects the connection between the application and the cache. The application can retrieve the AUTH token from Secrets Manager and present it when connecting.

https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/encryption.html
https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/auth.html
https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/in-transit-encryption.html

✗ **C** Create an SSL certificate, store it in AWS Secrets Manager, create a new cluster with encryption in transit, and update the application to retrieve and use the certificate for authentication.
Encryption in transit uses TLS, but an SSL certificate is not the Redis AUTH token used to authenticate clients. This option therefore does not satisfy the requirement to require user authentication.

https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/auth.html
https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/in-transit-encryption.html

✗ **D** Create an SSL certificate, store it in AWS Systems Manager Parameter Store as an encrypted advanced parameter, update the existing cluster to configure encryption in transit, and update the application to retrieve and use the certificate for authentication.
Enabling encryption in transit addresses connection security, but a client-provided SSL certificate is not the Redis AUTH credential required for user authentication. The application must authenticate with an AUTH token instead.

https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/auth.html
https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/in-transit-encryption.html

Question 66

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Improve the reliability of existing solutions.

A company uses an AWS CodeCommit repository The company must store a backup copy of the data that is in the repository in a second AWS Region Which solution will meet these requirements?

A. Configure AWS Elastic Disaster Recovery to replicate the CodeCommit repository data to the second Region

B. Use AWS Backup to back up the CodeCommit repository on an hourly schedule Create a cross-Region copy in the second Region

C. Create an Amazon EventBridge rule to invoke AWS CodeBuild when the company pushes code to the repository Use CodeBuild to clone the repository Create a zip file of the content Copy the file to an S3 bucket in the second Region

D. Create an AWS Step Functions workflow on an hourly schedule to take a snapshot of the CodeCommit repository Configure the workflow to copy the snapshot to an S3 bucket in the second Region

Answer: C. Create an Amazon EventBridge rule to invoke AWS CodeBuild when the company pushes code to the repository Use CodeBuild to clone the repository Create a zip file of the content Copy the file to an S3 bucket in the second Region

Explanation

Option C creates a backup copy by responding to repository push events, having CodeBuild clone the

Free demo · full exam: 681 questions

repository, and copying the resulting archive to an S3 bucket in the second Region. CodeCommit supports repository event notifications that can be used with EventBridge, enabling the workflow to run when code is pushed [9]. The other options rely on services or snapshot capabilities that do not provide the described backup of a CodeCommit repository.

https://docs.aws.amazon.com/codecommit/latest/userguide/monitoring-events.html

✗ **A** Configure AWS Elastic Disaster Recovery to replicate the CodeCommit repository data to the second Region
Elastic Disaster Recovery is designed to replicate server workloads, not the contents of a managed CodeCommit repository. It therefore does not provide the required repository backup.

https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html

✗ **B** Use AWS Backup to back up the CodeCommit repository on an hourly schedule Create a cross- Region copy in the second Region
This option assumes AWS Backup can create backups and cross-Region copies of CodeCommit repositories. It does not provide the described CodeCommit repository backup workflow.

✓ **C** Create an Amazon EventBridge rule to invoke AWS CodeBuild when the company pushes code to the repository Use CodeBuild to clone the repository Create a zip file of the content Copy the file to an S3 bucket in the second Region
A CodeCommit push event can be used to trigger an EventBridge rule [9]. CodeBuild can then clone the repository and create an archive, which can be copied to an S3 bucket in the second Region to retain a backup copy of the repository data.

https://docs.aws.amazon.com/codecommit/latest/userguide/monitoring-events.html

✗ **D** Create an AWS Step Functions workflow on an hourly schedule to take a snapshot of the CodeCommit repository Configure the workflow to copy the snapshot to an S3 bucket in the second Region
The proposed workflow depends on taking a snapshot of the CodeCommit repository, but the option does not specify a supported mechanism for creating such a snapshot. A schedule and cross-Region S3 destination do not by themselves provide the repository contents.

Dumpsgate

Question 67

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Determine a strategy to improve overall operational excellence.

A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.

The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output. The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.

A solutions architect needs to simplify this process to minimize disruption to users.

Which solution will meet these requirements with the LEAST operational overhead?

A. Directly modify the environment variables of the published Lambda function version. Use the SLATEST version to test image processing parameters.

B. Create an Amazon DynamoDB table to store the image processing parameters. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.

C. Directly code the image processing parameters within the Lambda function and remove the environment variables. Publish a new function version when the company updates the parameters.

D. Create a Lambda function alias. Modify the client application to use the function alias ARN. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Answer: D. Create a Lambda function alias. Modify the client application to use the function alias ARN. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Explanation

A Lambda alias provides a stable ARN for the application to invoke while allowing the company to change which published function version receives requests. The company can test updated environment variables in a new version and then repoint the alias, so it does not need to update the application for each release. This avoids disruptions from frequent client changes and requires less operational effort than adding a separate parameter store.

https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html
https://docs.aws.amazon.com/lambda/latest/dg/using-aliases.html

✗ **A** Directly modify the environment variables of the published Lambda function version. Use the $LATEST version to test image processing parameters.
A published Lambda version is an immutable snapshot, so its environment variables cannot be directly modified. Testing with $LATEST does not solve the client's need to invoke a new published version ARN after validation.

https://docs.aws.amazon.com/lambda/latest/dg/configuration-envvars.html?icmpid=docs_lambda_help
https://docs.aws.amazon.com/en_gb/lambda/latest/dg/configuration-versions.html

✗ **B** Create an Amazon DynamoDB table to store the image processing parameters. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.
This could let the company change parameters without publishing a new function version, but it requires

Free demo · full exam: 681 questions

creating and maintaining a DynamoDB table and adding retrieval logic to the function. The alias approach meets the requirement with less operational overhead and avoids application updates when a new version is ready.

✗ **C** Directly code the image processing parameters within the Lambda function and remove the environment variables. Publish a new function version when the company updates the parameters. This still requires publishing a new version whenever parameters change. Because the application invokes a version-specific ARN, it would also continue to need changes for each new version, causing the same user disruption.

https://docs.aws.amazon.com/en_gb/lambda/latest/dg/configuration-versions.html

✓ **D** Create a Lambda function alias. Modify the client application to use the function alias ARN. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing. The application can invoke the stable alias ARN instead of a specific version ARN. After testing and publishing a version with updated environment variables, the company can repoint the alias to that version without changing the application.

https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html
https://docs.aws.amazon.com/lambda/latest/dg/using-aliases.html

Question 68

Single choice · Domain: D3 Continuous Improvement for Existing Solutions › Recommend improvements to an existing solution to reduce costs.

A company is using an Amazon ECS cluster to run a data-processing application. Different business groups share ECS services in the ECS cluster. The ECS cluster runs on Amazon EC2 instances. ECS cluster auto scaling is enabled.

The company needs to assign EC2 costs of ECS tasks to the appropriate business groups.

Which solution will meet this requirement with the LEAST operational overhead?

A. Create a cost allocation tag on the EC2 Auto Scaling group to indicate the business group. Use AWS Cost Explorer to assign EC2 costs to the appropriate business group.

B. Enable split cost allocation data in AWS Cost Explorer. Create an AWS Cost and Usage Report that uses tags to assign EC2 costs to the appropriate business group.

C. Create a separate ECS cluster for each business group. Use AWS Cost Explorer to assign EC2 costs to the appropriate business group.

D. Create an AWS cost category for each business group. Define split charge rules for the ECS cluster for the business groups. Create an AWS Cost and Usage Report.

Answer: B. Enable split cost allocation data in AWS Cost Explorer. Create an AWS Cost and Usage Report that uses tags to assign EC2 costs to the appropriate business group.

Explanation

Split cost allocation data provides a way to attribute shared EC2 instance costs to ECS tasks, rather than assigning the entire instance cost to a business group. With the relevant ECS task tags and split cost allocation data included in a Cost and Usage Report, the company can use those records to allocate costs to the appropriate groups without creating separate clusters. This meets the attribution requirement with less operational overhead than restructuring the ECS environment.

https://docs.aws.amazon.com/cur/latest/userguide/split-cost-allocation-data.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/cur/latest/userguide/example-split-cost-allocation-data.html
https://docs.aws.amazon.com/cur/latest/userguide/enabling-split-cost-allocation-data.html
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-using-tags.html

✗ **A** Create a cost allocation tag on the EC2 Auto Scaling group to indicate the business group. Use AWS Cost Explorer to assign EC2 costs to the appropriate business group.
An Auto Scaling group tag can identify the group associated with the EC2 capacity, but it does not divide shared instance costs among the ECS tasks running on that capacity. Because different business groups share ECS services in the cluster, this approach cannot accurately attribute each task's share of the EC2 costs.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-tagging.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/add-tags.html

✓ **B** Enable split cost allocation data in AWS Cost Explorer. Create an AWS Cost and Usage Report that uses tags to assign EC2 costs to the appropriate business group.
Split cost allocation data provides task-level cost allocation for ECS tasks running on shared EC2 instances, instead of treating the instance cost as belonging to just one group. A Cost and Usage Report can include this split data and relevant task tags so costs can be assigned to the appropriate business groups without creating separate clusters.

https://docs.aws.amazon.com/cur/latest/userguide/split-cost-allocation-data.html
https://docs.aws.amazon.com/cur/latest/userguide/example-split-cost-allocation-data.html
https://docs.aws.amazon.com/cur/latest/userguide/enabling-split-cost-allocation-data.html
https://docs.aws.amazon.com/cur/latest/userguide/split-line-item-columns.html
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-using-tags.html

✗ **C** Create a separate ECS cluster for each business group. Use AWS Cost Explorer to assign EC2 costs to the appropriate business group.
Separate clusters could make it easier to associate capacity with a business group, but they require restructuring and managing additional clusters. That is more operational overhead than enabling split cost allocation data to attribute costs within the existing shared cluster.

https://docs.aws.amazon.com/cur/latest/userguide/split-cost-allocation-data.html
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/cluster-auto-scaling.html

✗ **D** Create an AWS cost category for each business group. Define split charge rules for the ECS cluster for the business groups. Create an AWS Cost and Usage Report.
Cost category split charge rules distribute costs according to configured allocation rules; they do not provide the ECS task-level cost attribution needed to apportion shared EC2 costs based on the tasks using them. This adds rule configuration while failing to use split cost allocation data for the task-level breakdown.

https://docs.aws.amazon.com/cur/latest/userguide/split-cost-allocation-data.html
https://docs.aws.amazon.com/aws-cost-management/latest/APIReference/API_CostCategorySplitChargeRule.html

Free demo · full exam: 681 questions

Question 69

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Select existing workloads and processes for potential migration.

A large company runs an AWS migration assessment. Now, the company must select its first workload to migrate to AWS. The company has the following goals:

- Build cloud migration skills.
- Minimize business risk.
- Move quickly to prepare for an on-premises data center contract expiration in less than 6 months.

Which migration approach aligns with AWS guidance for selecting and migrating a first workload?

A. Select a tightly coupled claims-processing system workload. Re-architect into microservices on Amazon ECS and Amazon RDS to modernize early.

B. Select an internal marketing portal workload. Rehost by using AWS Application Migration Service. Establish ongoing patch management by using AWS Systems Manager Patch Manager.

C. Select a low-latency trading system workload. Lift and shift by using AWS Direct Connect with custom packet serialization and proprietary vendor drivers.

D. Select a compliance reporting batch workload. Replatform the database to Amazon Aurora PostgreSQL by using AWS DMS with schema conversion.

Answer: B. Select an internal marketing portal workload. Rehost by using AWS Application Migration Service. Establish ongoing patch management by using AWS Systems Manager Patch Manager.

Explanation

The internal marketing portal is a suitable first workload because rehosting it with AWS Application Migration Service supports a relatively direct move, helping the company make progress before its data center contract expires. Choosing a less critical portal rather than a tightly coupled or latency-sensitive system helps minimize business risk while the team builds migration experience. AWS Systems Manager Patch Manager can support ongoing patch management for the migrated workload.

https://docs.aws.amazon.com/us_en/systems-manager/latest/userguide/patch-manager.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/prioritization-and-migration-strategy.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/aws-application-design-and-migration-strategy.html

✗ **A** Select a tightly coupled claims-processing system workload. Re-architect into microservices on Amazon ECS and Amazon RDS to modernize early.
Re-architecting a tightly coupled claims-processing system is a substantial change to a business workload, making it a risky choice for the first migration. It is also likely to require more time and specialized effort than a direct move, which conflicts with the less-than-six-month deadline.

https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/prioritization-and-migration-strategy.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/aws-application-design-and-migration-strategy.html

✓ **B** Select an internal marketing portal workload. Rehost by using AWS Application Migration Service. Establish ongoing patch management by using AWS Systems Manager Patch Manager.
Rehosting an internal portal is a relatively direct migration approach that can help the company move quickly while limiting the risk of changing a tightly coupled or business-critical system. Using the

Free demo · full exam: 681 questions

migration as an initial workload also gives the team practical migration experience. Patch Manager supports establishing ongoing patch management for the migrated systems.

https://docs.aws.amazon.com/us_en/systems-manager/latest/userguide/patch-manager.html
https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-installing-patches.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/prioritization-and-migration-strategy.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/aws-application-design-and-migration-strategy.html

✗ **C** Select a low-latency trading system workload. Lift and shift by using AWS Direct Connect with custom packet serialization and proprietary vendor drivers.
A low-latency trading system is a poor choice for minimizing risk on a first migration because its performance requirements can make migration and validation more demanding. The custom serialization and proprietary drivers add complexity, which is at odds with moving quickly and building skills through a manageable first workload.

https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/prioritization-and-migration-strategy.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/aws-application-design-and-migration-strategy.html

✗ **D** Select a compliance reporting batch workload. Replatform the database to Amazon Aurora PostgreSQL by using AWS DMS with schema conversion.
Replatforming the database and converting its schema introduces additional change and validation work compared with rehosting. That makes this a less suitable first workload when the company needs to limit risk, gain migration experience, and move quickly before the contract expires.

https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/prioritization-and-migration-strategy.html
https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/aws-application-design-and-migration-strategy.html

Dumpsgate

Question 70

Multiple choice · Domain: D4 Accelerate Workload Migration and Modernization › Determine the right migration approach for existing workloads.

A company wants to back up its on-premises intranet application to the AWS Cloud. The company wants to use AWS Elastic Disaster Recovery for the backup solution.

The company requires that replication traffic must travel across a private dedicated connection. The application must not be accessible from the public internet. The solution cannot consume more than 75% of the available bandwidth between the on-premises network and AWS.

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

A. Create a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway.

B. Create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway.

C. Create an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network. Configure Quality of Service (QoS) to ensure that the backup does not exceed 75% of the available bandwidth.

D. Create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network.

E. Configure the replication servers to use private IP addresses for data replication. Configure throttling on the replication instances so that the Elastic Disaster Recovery server bandwidth does not exceed 75% of the available bandwidth.

F. Configure the recovery instance's private IP address to match the source server's private IP address. Configure throttling on the recovery instances so that the total bandwidth the recovery instances consume does not exceed 75% of the available bandwidth.

Answer

- A. **Create a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway.**
- D. **Create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network.**
- E. **Configure the replication servers to use private IP addresses for data replication. Configure throttling on the replication instances so that the Elastic Disaster Recovery server bandwidth does not exceed 75% of the available bandwidth.**

Explanation

Use private subnets for the staging resources, with NAT gateways for outbound connectivity without making the application publicly accessible. AWS Direct Connect provides the required private dedicated connection, and a Direct Connect gateway can connect that path to the target AWS network. Configure Elastic Disaster Recovery replication-server throttling and private IP replication so replication uses no more than 75% of the available bandwidth.

https://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html
https://docs.aws.amazon.com/directconnect/latest/UserGuide/virtualgateways.html
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html
https://docs.aws.amazon.com/drs/latest/userguide/replication-settings.html

✓ **A** Create a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway. Private subnets keep the staging resources off the public internet, while NAT gateways allow outbound

Free demo · full exam: 681 questions

connectivity without allowing unsolicited inbound connections. The virtual private gateway provides the AWS-side attachment for the private connection; using two subnets and NAT gateways supports a multi-subnet design.

https://docs.aws.amazon.com/directconnect/latest/UserGuide/virtualgateways.html
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html
https://docs.aws.amazon.com/drs/latest/userguide/Network-Requirements.html

✗ **B** Create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway.
This design uses public subnets and an internet gateway, which conflicts with the requirement that the application not be accessible from the public internet. A virtual private gateway alone does not make those subnets private.

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Internet_Gateway.html

✗ **C** Create an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network. Configure Quality of Service (QoS) to ensure that the backup does not exceed 75% of the available bandwidth.
The requirement is for a private dedicated connection, whereas this choice specifies a Site-to-Site VPN rather than AWS Direct Connect. It also relies on QoS instead of the Elastic Disaster Recovery replication-server throttling called for to limit the replication bandwidth.

https://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html
https://docs.aws.amazon.com/drs/latest/userguide/replication-settings.html

✓ **D** Create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network.
AWS Direct Connect provides a dedicated network connection from the on-premises network to AWS, meeting the private dedicated connection requirement. A Direct Connect gateway can connect the Direct Connect path to the target AWS network through its virtual private gateway.

https://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html
https://docs.aws.amazon.com/directconnect/latest/UserGuide/virtualgateways.html
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html

✓ **E** Configure the replication servers to use private IP addresses for data replication. Configure throttling on the replication instances so that the Elastic Disaster Recovery server bandwidth does not exceed 75% of the available bandwidth.
Using private IP addresses keeps replication traffic on the private network path rather than requiring public IP connectivity. Throttling the replication servers limits the bandwidth consumed by replication to the required maximum of 75%.

https://docs.aws.amazon.com/drs/latest/userguide/replication-settings.html
https://docs.aws.amazon.com/drs/latest/userguide/Network-Requirements.html

✗ **F** Configure the recovery instance’s private IP address to match the source server's private IP address. Configure throttling on the recovery instances so that the total bandwidth the recovery instances consume does not exceed 75% of the available bandwidth.
Recovery instances are not the Elastic Disaster Recovery replication servers that transfer source data during replication. Matching a recovery instance's IP address and throttling recovery instances therefore do not ensure that replication traffic stays within the 75% bandwidth limit.

https://docs.aws.amazon.com/drs/latest/userguide/replication-settings.html

Free demo · full exam: 681 questions

Question 71

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Determine the right migration approach for existing workloads.

An online gaming company needs to rehost its gaming platform on AWS. The company's gaming application requires high performance computing (HPC) processing. The application has a leaderboard that updates frequently. The company uses an Amazon EC2 compute optimized instance that runs Ubuntu to host a Node.js application to display the game. The company tracks the state of the application in an on-premises Redis instance.

The company needs a migration strategy that optimizes application performance.

Which solution will meet these requirements?

A. Create an Auto Scaling group of memory optimized Amazon EC2 Spot Instances behind an Application Load Balancer. Use an Amazon ElastiCache (Redis OSS) cluster to maintain the leaderboard.

B. Create an Auto Scaling group of compute optimized Amazon EC2 Spot Instances behind an Application Load Balancer. Use an Amazon OpenSearch Service cluster to maintain the leaderboard.

C. Create an Auto Scaling group of compute optimized Amazon EC2 On-Demand Instances behind an Application Load Balancer. Use an Amazon ElastiCache (Redis OSS) cluster to maintain the leaderboard.

D. Create an Auto Scaling group of memory optimized Amazon EC2 On-Demand Instances behind an Application Load Balancer. Use an Amazon DynamoDB table to maintain the leaderboard.

Answer: C. Create an Auto Scaling group of compute optimized Amazon EC2 On-Demand Instances behind an Application Load Balancer. Use an Amazon ElastiCache (Redis OSS) cluster to maintain the leaderboard.

Explanation

Option C keeps the application on compute optimized instances, which fit the company's HPC processing needs and its existing compute optimized deployment. On-Demand Instances avoid the interruption risk associated with Spot capacity for a frequently used gaming application. ElastiCache for Redis OSS provides a managed Redis-compatible destination for the application's frequently updated leaderboard state.

https://docs.aws.amazon.com/ec2/latest/instancetypes/co.html
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-interruptions.html

✗ **A** Create an Auto Scaling group of memory optimized Amazon EC2 Spot Instances behind an Application Load Balancer. Use an Amazon ElastiCache (Redis OSS) cluster to maintain the leaderboard. ElastiCache for Redis OSS suits the leaderboard, but memory optimized instances do not match the application's stated HPC processing need as well as compute optimized instances. Spot capacity can also be interrupted, which is a poor fit for the primary gaming application when consistent availability is important.

https://docs.aws.amazon.com/ec2/latest/instancetypes/co.html
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-interruptions.html

✗ **B** Create an Auto Scaling group of compute optimized Amazon EC2 Spot Instances behind an Application Load Balancer. Use an Amazon OpenSearch Service cluster to maintain the leaderboard. Compute optimized instances fit the HPC processing requirement, but Spot Instances can be interrupted. OpenSearch is not the Redis-compatible in-memory store that matches the company's existing Redis-based leaderboard state and its frequent updates.

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/ec2/latest/instancetypes/co.html
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-interruptions.html

✓ **C** Create an Auto Scaling group of compute optimized Amazon EC2 On-Demand Instances behind an Application Load Balancer. Use an Amazon ElastiCache (Redis OSS) cluster to maintain the leaderboard. Compute optimized instances are suited to the application's compute-intensive HPC processing. On-Demand capacity avoids Spot interruption risk, while ElastiCache for Redis OSS provides a managed destination for the on-premises Redis state used by the frequently updated leaderboard.

https://docs.aws.amazon.com/ec2/latest/instancetypes/co.html
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-interruptions.html

✗ **D** Create an Auto Scaling group of memory optimized Amazon EC2 On-Demand Instances behind an Application Load Balancer. Use an Amazon DynamoDB table to maintain the leaderboard.
On-Demand capacity avoids Spot interruptions, but memory optimized instances are not the best match for the stated HPC processing requirement compared with compute optimized instances. DynamoDB also replaces the existing Redis-based state store rather than preserving the Redis-compatible approach suited to the frequently updated leaderboard.

https://docs.aws.amazon.com/ec2/latest/instancetypes/co.html

Question 72

Multiple choice · Domain: D4 Accelerate Workload Migration and Modernization › Determine a new architecture for existing workloads.

A company has a platform that consists of an on-premises single-node Kubernetes cluster. The cluster uses Windows-based and Linux-based file storage as shared storage. The company wants to migrate its platform to AWS.

The cluster has multiple containerized applications that are continuously scanned for operating system vulnerabilities and programming language package vulnerabilities. The company needs a highly available solution that supports monitoring, logging, and multiprotocol shared storage.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Select TWO.)

A. Replace the existing Kubernetes cluster with a single-node Amazon EKS cluster that is managed by AWS. Install Prometheus and Grafana on the EKS cluster by using public Docker images. Store the application Docker images in Amazon ECR. Enable ECR enhanced scanning.

B. Replace the existing Kubernetes cluster with a multi-node Amazon EKS cluster that is managed by AWS. Set up Amazon Managed Service for Prometheus and Amazon Managed Grafana for monitoring. Store the application Docker images in Amazon ECR. Enable ECR enhanced scanning.

C. Replace the existing Kubernetes cluster with a multi-node Amazon EKS cluster that is managed by AWS. Install open source versions of Prometheus and Grafana on Amazon EC2 instances. Store the application Docker images in Amazon ECR. Enable ECR enhanced scanning.

D. Provision an Amazon FSx for Windows File Server file share and an Amazon EFS file system. Mount the file share and the file system on the containers launched on the Amazon EKS cluster as persistent volumes.

E. Provision an Amazon FSx for NetApp ONTAP file system. Mount the file system on the containers launched on the Amazon EKS cluster as a persistent volume.

Answer

- B. **Replace the existing Kubernetes cluster with a multi-node Amazon EKS cluster that is managed by AWS. Set up Amazon Managed Service for Prometheus and Amazon Managed Grafana for monitoring. Store the application Docker images in Amazon ECR. Enable ECR enhanced scanning.**
- E. **Provision an Amazon FSx for NetApp ONTAP file system. Mount the file system on the containers launched on the Amazon EKS cluster as a persistent volume.**

Explanation

B and E provide the requested highly available, lower-overhead platform: a multi-node Amazon EKS cluster avoids relying on a single node, while managed Prometheus and Grafana reduce the work of operating monitoring components. ECR enhanced scanning addresses the application image vulnerability-scanning requirement. FSx for NetApp ONTAP provides shared storage that supports both Windows and Linux access protocols, avoiding the need to operate separate storage systems for the two environments.

https://docs.aws.amazon.com/prometheus/latest/userguide/what-is-Amazon-Managed-Service-Prometheus.html
https://docs.aws.amazon.com/eks/latest/userguide/disaster-recovery-resiliency.html
https://docs.aws.amazon.com/eks/latest/best-practices/data-plane.html
https://docs.aws.amazon.com/grafana/latest/userguide/AMG_solutions.html

✗ **A** Replace the existing Kubernetes cluster with a single-node Amazon EKS cluster that is managed by AWS. Install Prometheus and Grafana on the EKS cluster by using public Docker images. Store the application Docker images in Amazon ECR. Enable ECR enhanced scanning.
A single-node cluster does not meet the high-availability requirement. Installing and maintaining Prometheus and Grafana on the cluster also creates more operational work than using the managed services in option B. ECR enhanced scanning is useful for image vulnerability scanning, but it does not correct those shortcomings.

https://docs.aws.amazon.com/prometheus/latest/userguide/what-is-Amazon-Managed-Service-Prometheus.html
https://docs.aws.amazon.com/eks/latest/userguide/disaster-recovery-resiliency.html
https://docs.aws.amazon.com/eks/latest/best-practices/data-plane.html

✓ **B** Replace the existing Kubernetes cluster with a multi-node Amazon EKS cluster that is managed by AWS. Set up Amazon Managed Service for Prometheus and Amazon Managed Grafana for monitoring. Store the application Docker images in Amazon ECR. Enable ECR enhanced scanning.
A multi-node EKS cluster supports a more highly available deployment than the existing single-node cluster. Amazon Managed Service for Prometheus and Amazon Managed Grafana reduce the operational burden of running those monitoring components, and ECR enhanced scanning supports scanning images for operating system and programming-language package vulnerabilities.

https://docs.aws.amazon.com/prometheus/latest/userguide/what-is-Amazon-Managed-Service-Prometheus.html
https://docs.aws.amazon.com/eks/latest/userguide/disaster-recovery-resiliency.html
https://docs.aws.amazon.com/eks/latest/best-practices/data-plane.html
https://docs.aws.amazon.com/grafana/latest/userguide/AMG_solutions.html

✗ **C** Replace the existing Kubernetes cluster with a multi-node Amazon EKS cluster that is managed by AWS. Install open source versions of Prometheus and Grafana on Amazon EC2 instances. Store the application Docker images in Amazon ECR. Enable ECR enhanced scanning.
The multi-node EKS cluster and ECR scanning address important requirements, but hosting Prometheus and Grafana on EC2 means the company must operate and maintain those instances and software. That adds operational overhead compared with the managed monitoring services in B.

https://docs.aws.amazon.com/prometheus/latest/userguide/what-is-Amazon-Managed-Service-Prometheus.html
https://docs.aws.amazon.com/eks/latest/userguide/disaster-recovery-resiliency.html
https://docs.aws.amazon.com/eks/latest/best-practices/data-plane.html

✗ **D** Provision an Amazon FSx for Windows File Server file share and an Amazon EFS file system. Mount the file share and the file system on the containers launched on the Amazon EKS cluster as persistent

Free demo · full exam: 681 questions

volumes.
This provisions separate storage systems for Windows-oriented and Linux-oriented access rather than one multiprotocol shared-storage solution. FSx for NetApp ONTAP in E provides a single filesystem that can serve both Windows and Linux clients through supported protocols.

✓ **E** Provision an Amazon FSx for NetApp ONTAP file system. Mount the file system on the containers launched on the Amazon EKS cluster as a persistent volume.
FSx for NetApp ONTAP supports multiprotocol access, allowing the shared-storage requirement for Windows-based and Linux-based workloads to be met with one file system. Using it as persistent storage for the EKS containers avoids operating separate Windows and Linux storage services.

Question 73

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Select existing workloads and processes for potential migration.

A company is migrating to AWS and needs to inventory physical and virtual servers, apps, and database relationships to properly rightsize and plan migration.

A. Use Migration Evaluator with Agentless Collector.
B. Use Migration Hub with Discovery Agent and Strategy Recommendations.
C. Use Migration Hub with Agentless Collector and Migration Service.
D. Use Migration Hub import tool.

Answer: B. Use Migration Hub with Discovery Agent and Strategy Recommendations.

Explanation

Migration Hub with the Discovery Agent and Strategy Recommendations best matches the need to inventory servers and applications, understand their relationships, and use that information to plan migration and rightsize. The Discovery Agent gathers detailed server and dependency data, while Strategy Recommendations analyzes discovery data to provide migration and modernization guidance. Together, they address both discovery and planning rather than only importing data or carrying out a migration.

https://docs.aws.amazon.com/migrationhub-strategy/latest/userguide/application-data-collector.html
https://docs.aws.amazon.com/migrationhub-strategy/latest/userguide/getting-started-prerequisites.html
https://docs.aws.amazon.com/migrationhub-strategy/latest/userguide/getting-started-get-recommendations.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html

✗ **A** Use Migration Evaluator with Agentless Collector.
This is not the choice that pairs application discovery with Migration Hub Strategy Recommendations. The scenario calls for inventory and relationship information that can be analyzed to plan migration and rightsize; Strategy Recommendations provides that planning analysis using discovery data.

https://docs.aws.amazon.com/migrationhub-strategy/latest/userguide/getting-started-get-recommendations.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/agentless-collector-qs-prerequisites.html

✓ **B** Use Migration Hub with Discovery Agent and Strategy Recommendations.
The Discovery Agent collects detailed server information, including data useful for identifying application dependencies. Strategy Recommendations analyzes discovery data to help assess migration and modernization options, fitting the need to inventory the environment and plan migration and rightsize.

https://docs.aws.amazon.com/migrationhub-strategy/latest/userguide/application-data-collector.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/migrationhub-strategy/latest/userguide/getting-started-prerequisites.html
https://docs.aws.amazon.com/migrationhub-strategy/latest/userguide/getting-started-get-recommendations.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html

✗ **C** Use Migration Hub with Agentless Collector and Migration Service.
The Agentless Collector can collect inventory data, but this option pairs it with a migration service rather than Strategy Recommendations for analyzing the inventory and planning the migration. The scenario needs assessment and recommendations, not just data collection and migration execution.

https://docs.aws.amazon.com/migrationhub-strategy/latest/userguide/getting-started-get-recommendations.html
https://docs.aws.amazon.com/migrationhub/latest/ug/gs-new-user-migration.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/agentless-collector-gs-prerequisites.html

✗ **D** Use Migration Hub import tool.
The import tool is for bringing existing discovery data into the service; it does not itself discover the company's servers, applications, and relationships. It also does not provide the Strategy Recommendations analysis needed for migration planning and rightsizing.

https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-import.html
https://docs.aws.amazon.com/migrationhub-strategy/latest/userguide/getting-started-get-recommendations.html

Question 74

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Select existing workloads and processes for potential migration.

A company wants to migrate to AWS. The company is running thousands of VMs in a VMware ESXi environment. The company has no configuration management database and has little Knowledge about the utilization of the VMware portfolio.

A solutions architect must provide the company with an accurate inventory so that the company can plan for a cost-effective migration.

Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Systems Manager Patch Manager to deploy Migration Evaluator to each VM. Review the collected data in Amazon QuickSight. Identify servers that have high utilization. Remove the servers that have high utilization from the migration list. Import the data to AWS Migration Hub.

B. Export the VMware portfolio to a csv file. Check the disk utilization for each server. Remove servers that have high utilization. Export the data to AWS Application Migration Service. Use AWS Server Migration Service (AWS SMS) to migrate the remaining servers.

C. Deploy the Migration Evaluator agentless collector to the ESXi hypervisor. Review the collected data in Migration Evaluator. Identify inactive servers. Remove the inactive servers from the migration list. Import the data to AWS Migration Hub.

D. Deploy the AWS Application Migration Service Agent to each VM. When the data is collected, use Amazon Redshift to import and analyze the data. Use Amazon QuickSight for data visualization.

Answer: C. Deploy the Migration Evaluator agentless collector to the ESXi hypervisor. Review the collected data in Migration Evaluator. Identify inactive servers. Remove the inactive servers from the migration list. Import the data to AWS Migration Hub.

Explanation

The agentless collector can gather inventory and utilization information from the VMware ESXi

Free demo · full exam: 681 questions

environment without installing software on thousands of individual VMs. Reviewing the results in Migration Evaluator helps the company identify inactive servers to exclude from migration planning, supporting a more cost-effective migration. This approach avoids the manual work or per-VM deployment effort in the other options; the resulting discovery data can also be imported into AWS Migration Hub.

https://docs.aws.amazon.com/application-discovery/latest/userguide/agentless-collector-dashboard.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/agentless-collector-qs-view-collected-data.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-import.html

✗ **A** Use AWS Systems Manager Patch Manager to deploy Migration Evaluator to each VM. Review the collected data in Amazon QuickSight. Identify servers that have high utilization. Remove the servers that have high utilization from the migration list. Import the data to AWS Migration Hub.
This proposes deploying a tool to each VM, which adds operational effort compared with collecting data from the ESXi environment agentlessly. It also removes highly utilized servers, rather than identifying inactive servers to exclude from the migration plan; high utilization alone does not make a server an unsuitable migration candidate.

✗ **B** Export the VMware portfolio to a csv file. Check the disk utilization for each server. Remove servers that have high utilization. Export the data to AWS Application Migration Service. Use AWS Server Migration Service (AWS SMS) to migrate the remaining servers.
Manually exporting and checking servers is labor-intensive and checks disk utilization rather than providing the requested broader inventory and utilization assessment. It also excludes highly utilized servers, which does not identify inactive servers for cost-effective migration planning.

✓ **C** Deploy the Migration Evaluator agentless collector to the ESXi hypervisor. Review the collected data in Migration Evaluator. Identify inactive servers. Remove the inactive servers from the migration list. Import the data to AWS Migration Hub.
The agentless collector gathers information from the VMware environment without requiring installation on each VM, reducing operational overhead at this scale. Migration Evaluator can be used to review the collected data and identify inactive servers to exclude from the migration plan; discovery data can be imported into Migration Hub.

https://docs.aws.amazon.com/application-discovery/latest/userguide/agentless-collector-dashboard.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/agentless-collector-qs-view-collected-data.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-import.html

✗ **D** Deploy the AWS Application Migration Service Agent to each VM. When the data is collected, use Amazon Redshift to import and analyze the data. Use Amazon QuickSight for data visualization.
Installing an agent on every VM adds deployment and management effort, and this option uses a migration replication agent rather than the agentless assessment approach suited to building an inventory. Importing and analyzing the data in Redshift is also unnecessary overhead for the stated inventory and migration-planning need.

Question 75

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Select existing workloads and processes for potential migration.

A company is planning to migrate an on-premises data center to AWS. The company currently hosts the data center on Linux-based VMware VMs. A solutions architect must collect information about network dependencies between the VMs. The information must be in the form of a diagram that details host IP addresses, hostnames, and network connection information.

Which solution will meet these requirements?

A. Use AWS Application Discovery Service. Select an AWS Migration Hub home AWS Region. Install the AWS Application Discovery Agent on the on-premises servers for data collection. Grant permissions to Application Discovery Service to use the Migration Hub network diagrams.

B. Use the AWS Application Discovery Service Agentless Collector for server data collection. Export the network diagrams from the AWS Migration Hub in .png format.

C. Install the AWS Application Migration Service agent on the on-premises servers for data collection. Use AWS Migration Hub data in Workload Discovery on AWS to generate network diagrams.

D. Install the AWS Application Migration Service agent on the on-premises servers for data collection. Export data from AWS Migration Hub in .csv format into an Amazon CloudWatch dashboard to generate network diagrams.

Answer: A. Use AWS Application Discovery Service. Select an AWS Migration Hub home AWS Region. Install the AWS Application Discovery Agent on the on-premises servers for data collection. Grant permissions to Application Discovery Service to use the Migration Hub network diagrams.

Explanation

AWS Migration Hub network diagrams use network connection data collected by the AWS Application Discovery Agent. The agent can collect the host and connection information needed to map dependencies, and Migration Hub must be configured with a home Region for discovery. Granting Application Discovery Service permission to use the diagrams completes the required setup. This provides a diagram with the host details and network connections requested.

https://docs.aws.amazon.com/migrationhub/latest/ug/network-diagram.html
https://docs.aws.amazon.com/migrationhub/latest/ug/gs-new-user-discovery.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html

✓ **A** Use AWS Application Discovery Service. Select an AWS Migration Hub home AWS Region. Install the AWS Application Discovery Agent on the on-premises servers for data collection. Grant permissions to Application Discovery Service to use the Migration Hub network diagrams.
This is the Migration Hub network-diagram workflow: the Application Discovery Agent collects server and network connection data, and Migration Hub uses that data to visualize dependencies. Selecting a home Region and granting the required permissions enables the service to use the network diagrams.

https://docs.aws.amazon.com/migrationhub/latest/ug/network-diagram.html
https://docs.aws.amazon.com/migrationhub/latest/ug/gs-new-user-discovery.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html

✗ **B** Use the AWS Application Discovery Service Agentless Collector for server data collection. Export the network diagrams from the AWS Migration Hub in .png format.
The described agentless collection approach does not provide the Application Discovery Agent-collected connection data used for Migration Hub network diagrams. The diagram feature depends on the discovery-agent network data, so exporting a diagram as an image does not address the missing

Free demo · full exam: 681 questions

collection method.

https://docs.aws.amazon.com/migrationhub/latest/ug/network-diagram.html
https://docs.aws.amazon.com/application-discovery/latest/userguide/agentless-collector-qs-network-data-collection.html

✗ **C** Install the AWS Application Migration Service agent on the on-premises servers for data collection. Use AWS Migration Hub data in Workload Discovery on AWS to generate network diagrams.
The AWS Application Migration Service agent is used for migration and replication, not for collecting the Application Discovery Service network-dependency data required by Migration Hub diagrams. Workload Discovery on AWS is not the specified way to generate these Migration Hub network diagrams.

https://docs.aws.amazon.com/migrationhub/latest/ug/network-diagram.html
https://docs.aws.amazon.com/pdfs/mgn/latest/ug/user-guide.pdf
https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html

✗ **D** Install the AWS Application Migration Service agent on the on-premises servers for data collection. Export data from AWS Migration Hub in .csv format into an Amazon CloudWatch dashboard to generate network diagrams.
The Application Migration Service agent does not collect the network-dependency data needed for Migration Hub network diagrams. Exporting CSV data to CloudWatch is also not the Migration Hub network-diagram workflow; the diagrams are generated from Application Discovery Service discovery data.

https://docs.aws.amazon.com/migrationhub/latest/ug/network-diagram.html
https://docs.aws.amazon.com/pdfs/mgn/latest/ug/user-guide.pdf
https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html

Question 76

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Determine the right migration approach for existing workloads.

A company wants to migrate an Amazon Aurora MySQL DB cluster from an existing AWS account to a new AWS account in the same AWS Region. Both accounts are members of the same organization in AWS Organizations.

The company must minimize database service interruption before the company performs DNS cutover to the new database.

Which migration strategy will meet this requirement?

A. Take a snapshot of the existing Aurora database. Share the snapshot with the new AWS account. Create an Aurora DB cluster in the new account from the snapshot.

B. Create an Aurora DB cluster in the new AWS account. Use AWS Database Migration Service (AWS DMS) to migrate data between the two Aurora DB clusters.

C. Use AWS Backup to share an Aurora database backup from the existing AWS account to the new AWS account. Create an Aurora DB cluster in the new AWS account from the snapshot.

D. Create an Aurora DB cluster in the new AWS account. Use AWS Application Migration Service to migrate data between the two Aurora DB clusters.

Answer: B. Create an Aurora DB cluster in the new AWS account. Use AWS Database Migration Service (AWS DMS) to migrate data between the two Aurora DB clusters.

Free demo · full exam: 681 questions

Dumpsgate

Question 77

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Select existing workloads and processes for potential migration.

A solutions architect needs to assess a newly acquired company's portfolio of applications and databases. The solutions architect must create a business case to migrate the portfolio to AWS. The newly acquired company runs applications in an on-premises data center. The data center is not well documented. The solutions architect cannot immediately determine how many applications and databases exist. Traffic for the applications is variable. Some applications are batch processes that run at the end of each month.

The solutions architect must gain a better understanding of the portfolio before a migration to AWS can begin.

Which solution will meet these requirements?

A. Use AWS Server Migration Service (AWS SMS) and AWS Database Migration Service (AWS DMS) to evaluate migration. Use AWS Service Catalog to understand application and database dependencies.

B. Use AWS Application Migration Service. Run agents on the on-premises infrastructure. Manage the agents by using AWS Migration Hub. Use AWS Storage Gateway to assess local storage needs and database dependencies.

C. Use Migration Evaluator to generate a list of servers. Build a report for a business case. Use AWS Migration Hub to view the portfolio. Use AWS Application Discovery Service to gain anunderstanding of application dependencies.

D. Use AWS Control Tower in the destination account to generate an application portfolio. Use AWS Server Migration Service (AWS SMS) to generate deeper reports and a business case. Use a landing zone for core accounts and resources.

Answer: C. Use Migration Evaluator to generate a list of servers. Build a report for a business case. Use AWS Migration Hub to view the portfolio. Use AWS Application Discovery Service to gain anunderstanding of application dependencies.

Explanation

Migration Evaluator can help inventory the on-premises environment and build a business case, addressing the unknown size of the portfolio and the need to assess it before migration. Migration Hub provides a portfolio view, while Application Discovery Service helps uncover application dependencies. Together, these services provide the discovery and assessment the architect needs without starting migration first.

https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html
https://docs.aws.amazon.com/migrationhub/
https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/directional-business-case.html
https://aws.amazon.com/migration-evaluator//

✗ **A** Use AWS Server Migration Service (AWS SMS) and AWS Database Migration Service (AWS DMS) to evaluate migration. Use AWS Service Catalog to understand application and database dependencies. SMS and DMS are migration services, not the right combination for discovering an undocumented portfolio and building its business case. AWS Service Catalog manages and provisions cataloged products; it does not discover application and database dependencies in the on-premises environment.

https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html
https://docs.aws.amazon.com/en_en/controltower/latest/userguide/service-catalog.html

✗ **B** Use AWS Application Migration Service. Run agents on the on-premises infrastructure. Manage the

Free demo · full exam: 681 questions

agents by using AWS Migration Hub. Use AWS Storage Gateway to assess local storage needs and database dependencies.
Application Migration Service is intended to migrate servers, rather than assess the portfolio and its dependencies before migration. Storage Gateway does not discover database dependencies, and Migration Hub is not used to manage Application Migration Service agents for this purpose.

https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html
https://docs.aws.amazon.com/migrationhub/

✓ **C** Use Migration Evaluator to generate a list of servers. Build a report for a business case. Use AWS Migration Hub to view the portfolio. Use AWS Application Discovery Service to gain an understanding of application dependencies.
Migration Evaluator supports assessing the environment and developing a business case, which is useful when the server count and workload patterns are not yet known. Migration Hub provides a portfolio view, and Application Discovery Service helps discover application dependencies before migration begins.

https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html
https://docs.aws.amazon.com/migrationhub/
https://docs.aws.amazon.com/prescriptive-guidance/latest/application-portfolio-assessment-guide/directional-business-case.html
https://aws.amazon.com/migration-evaluator//

✗ **D** Use AWS Control Tower in the destination account to generate an application portfolio. Use AWS Server Migration Service (AWS SMS) to generate deeper reports and a business case. Use a landing zone for core accounts and resources.
Control Tower helps set up and govern a multi-account AWS environment; it does not discover an on-premises application portfolio. A landing zone prepares AWS accounts and resources, while SMS is not the appropriate service for generating the discovery report and business case needed here.

https://docs.aws.amazon.com/en_en/controltower/latest/userguide/service-catalog.html
https://docs.aws.amazon.com/controlcatalog/latest/userguide/what-is-controlcatalog.html

Dumpsgate

Question 78

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Determine opportunities for modernization and enhancements.

A company wants to modernize a monolithic application in the company's data center and deploy the application on AWS. The monolithic application consists of an event broker in a central account and multiple microservices in individual AWS accounts. The event broker and the microservices are deployed on Amazon ECS clusters that use the Fargate launch type.

Multiple microservices need access to the same events from the event broker. The company wants to distribute events from the central event broker to each microservice across accounts.

Which solution will meet these requirements?

A. Create an Amazon SNS topic in the central account. Add a topic policy to allow other accounts to subscribe to the topic. Create an Amazon SQS queue in each individual AWS account. Subscribe the SQS queue to the SNS topic. Configure the microservices to read events from their own SQS queue.

B. Create a new Amazon EventBridge event bus in the central account with the required permissions. Add EventBridge rules filtered by service for each microservice. Invoke the rules to route events to other accounts.

C. Create a data stream in Amazon Kinesis Data Streams in the central account. Create an IAM policy to grant the necessary permissions to access the data stream. Set each of the microservices as an event source on the Kinesis stream. Configure the stream to invoke each microservice.

D. Create a new Amazon SQS queue as the event broker in the central account. Grant the required permissions. Configure each of the microservices to read messages from the central SQS queue.

Answer: B. Create a new Amazon EventBridge event bus in the central account with the required permissions. Add EventBridge rules filtered by service for each microservice. Invoke the rules to route events to other accounts.

Explanation

Use Amazon EventBridge to distribute events from the central account with rules that filter events for each microservice and route them across accounts. The central event bus must allow the relevant cross-account access, and the destination event buses must be permitted as targets. This supports routing selected events to each microservice's account rather than having all services compete to consume from one queue.

https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-event-bus-perms.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-service-cross-account.html
https://docs.aws.amazon.com/eventbridge/latest/APIReference/API_Target.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-targets.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-use-conditions.html

✗ **A** Create an Amazon SNS topic in the central account. Add a topic policy to allow other accounts to subscribe to the topic. Create an Amazon SQS queue in each individual AWS account. Subscribe the SQS queue to the SNS topic. Configure the microservices to read events from their own SQS queue.
SNS-to-SQS can provide fan-out in some designs, but this option creates a separate SNS/SQS distribution arrangement rather than using EventBridge rules to filter and route events from the central broker by service. EventBridge supports the cross-account event routing required by the intended solution.

https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-event-bus-perms.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-targets.html

Free demo · full exam: 681 questions

https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-use-conditions.html

✓ **B** Create a new Amazon EventBridge event bus in the central account with the required permissions. Add EventBridge rules filtered by service for each microservice. Invoke the rules to route events to other accounts.
EventBridge supports cross-account event-bus permissions and routing, so the central account can authorize delivery to other accounts. Rules can filter events for each microservice and route matching events to the appropriate account, allowing services to receive the events intended for them.

https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-event-bus-perms.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-service-cross-account.html
https://docs.aws.amazon.com/eventbridge/latest/APIReference/API_Target.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-targets.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-use-conditions.html

✗ **C** Create a data stream in Amazon Kinesis Data Streams in the central account. Create an IAM policy to grant the necessary permissions to access the data stream. Set each of the microservices as an event source on the Kinesis stream. Configure the stream to invoke each microservice.
A Kinesis stream does not directly invoke ECS services as event sources in the way this option describes. Kinesis event-source mappings are used with supported consumers such as Lambda, so this would not directly provide the specified routing to the Fargate microservices.

https://docs.aws.amazon.com/lambda/latest/dg/services-kinesis-create.html
https://docs.aws.amazon.com/lambda/latest/dg/with-kinesis-example.html

✗ **D** Create a new Amazon SQS queue as the event broker in the central account. Grant the required permissions. Configure each of the microservices to read messages from the central SQS queue.
When multiple microservices consume from the same SQS queue, they compete for messages; a message is not independently delivered to every service. Separate delivery to each service would require a fan-out design, whereas EventBridge rules can route selected events across accounts.

https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-targets.html

Dumpsgate

Question 79

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Determine opportunities for modernization and enhancements.

A company recently completed the migration from an on-premises data center to the AWS Cloud by using a replatforming strategy. One of the migrated servers is running a legacy Simple Mail Transfer Protocol (SMTP) service that a critical application relies upon. The application sends outbound email messages to the company's customers. The legacy SMTP server does not support TLS encryption and uses TCP port 25. The application can use SMTP only.

The company decides to use Amazon Simple Email Service (Amazon SES) and to decommission the legacy SMTP server. The company has created and validated the SES domain. The company has lifted the SES limits.

What should the company do to modify the application to send email messages from Amazon SES?

A. Configure the application to connect to Amazon SES by using TLS Wrapper. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Attach the IAM role to an Amazon EC2 instance.

B. Configure the application to connect to Amazon SES by using STARTTLS. Obtain Amazon SES SMTP credentials. Use the credentials to authenticate with Amazon SES.

C. Configure the application to use the SES API to send email messages. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Use the IAM role as a service role for Amazon SES.

D. Configure the application to use AWS SDKs to send email messages. Create an IAM user for Amazon SES. Generate API access keys. Use the access keys to authenticate with Amazon SES.

Answer: B. Configure the application to connect to Amazon SES by using STARTTLS. Obtain Amazon SES SMTP credentials. Use the credentials to authenticate with Amazon SES.

Explanation

Because the application can send email only through SMTP, it must connect to the Amazon SES SMTP interface rather than use the SES API or an SDK. Configure the application to use STARTTLS so the SMTP connection is encrypted, and authenticate with Amazon SES SMTP credentials. The SMTP credentials are distinct from using an IAM role or API access keys directly for SMTP authentication.

https://docs.aws.amazon.com/ses/latest/dg/send-email-concepts-credentials.html
https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html
https://docs.aws.amazon.com/ses/latest/dg/send-email-smtp.html
https://docs.aws.amazon.com/ses/latest/dg/smtp-credentials.html?icmpid=docs_ses_console

✗ **A** Configure the application to connect to Amazon SES by using TLS Wrapper. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Attach the IAM role to an Amazon EC2 instance. An IAM role attached to the EC2 instance does not provide the SMTP credentials the application needs to authenticate to the SES SMTP endpoint. This option also specifies TLS Wrapper rather than the STARTTLS configuration and SMTP credentials required by the correct approach.

https://docs.aws.amazon.com/ses/latest/dg/send-email-concepts-credentials.html
https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html
https://docs.aws.amazon.com/ses/latest/dg/smtp-credentials.html?icmpid=docs_ses_console

✓ **B** Configure the application to connect to Amazon SES by using STARTTLS. Obtain Amazon SES SMTP credentials. Use the credentials to authenticate with Amazon SES.
The application supports SMTP, so it can use the SES SMTP endpoint without switching to an API-based sending method. STARTTLS provides encryption for the SMTP connection, and SES SMTP credentials are

Free demo · full exam: 681 questions

used to authenticate the application.

https://docs.aws.amazon.com/ses/latest/dg/send-email-concepts-credentials.html
https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html
https://docs.aws.amazon.com/ses/latest/dg/send-email-smtp.html
https://docs.aws.amazon.com/ses/latest/dg/smtp-credentials.html?icmpid=docs_ses_console

✗ **C** Configure the application to use the SES API to send email messages. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Use the IAM role as a service role for Amazon SES.
The application can use SMTP only, so it cannot be changed to send messages through the SES API as this option requires. The IAM role and API permissions do not make the API approach compatible with the application's SMTP-only capability.

https://docs.aws.amazon.com/en_en/ses/latest/dg/control-user-access.html
https://docs.aws.amazon.com/ses/latest/APIReference-V2/API_SendEmail.html

✗ **D** Configure the application to use AWS SDKs to send email messages. Create an IAM user for Amazon SES. Generate API access keys. Use the access keys to authenticate with Amazon SES.
Using an AWS SDK to send email uses the SES API, which does not meet the requirement that the application can use SMTP only. IAM user API access keys are not the SES SMTP credentials the application needs for SMTP authentication.

https://docs.aws.amazon.com/ses/latest/dg/send-email-concepts-credentials.html
https://docs.aws.amazon.com/ses/latest/dg/send-email-smtp.html
https://docs.aws.amazon.com/ses/latest/dg/smtp-credentials.html?icmpid=docs_ses_console
https://docs.aws.amazon.com/ses/latest/APIReference-V2/API_SendEmail.html

Dumpsgate

Question 80

Single choice · Domain: D4 Accelerate Workload Migration and Modernization › Select existing workloads and processes for potential migration.

A company wants to migrate its website to AWS. The website uses containers that are deployed in an on-premises, self-managed Kubernetes cluster. All data for the website is stored in an on-premises PostgreSQL database.

The company has decided to migrate the on-premises Kubernetes cluster to an Amazon EKS cluster.

The EKS cluster will use EKS managed node groups with a static number of nodes. The company will also migrate the on-premises database to an Amazon RDS for PostgreSQL database.

A solutions architect needs to estimate the total cost of ownership (TCO) for this workload before the migration.

Which solution will provide the required TCO information?

A. Request access to Migration Evaluator. Run the Migration Evaluator Collector and import the data. Configure a scenario. Export a Quick Insights report from Migration Evaluator.

B. Launch AWS DMS for the on-premises database. Generate an assessment report. Create an estimate in AWS Pricing Calculator for the costs of the EKS migration.

C. Initialize AWS Application Migration Service. Add the on-premises servers as source servers. Launch a test instance. Output a TCO report from Application Migration Service.

D. Access the AWS Cloud Economics Center webpage to assess the AWS Cloud Value Framework. Create an AWS Cost and Usage report from the Cloud Value Framework.

Answer: A. Request access to Migration Evaluator. Run the Migration Evaluator Collector and import the data. Configure a scenario. Export a Quick Insights report from Migration Evaluator.

Free demo · full exam: 681 questions

Dumpsgate