

DUMPS ARENA

**AWS Certified Solutions Architect -
Professional**

Amazon AWS SAP-C02

Version Demo

Total Demo Questions: 134

Total Premium Questions: 1343

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Design Solutions for Organizational Complexity	299
Topic 2, Design for New Solutions	425
Topic 3, Continuous Improvement for Existing Solutions	389
Topic 4, Accelerate Workload Migration and Modernization	230
Total	1343

QUESTION NO: 1

A company is running an application in the AWS Cloud. The application uses AWS Lambda functions and Amazon Elastic Container Service (Amazon ECS) containers that run with AWS Fargate technology as its primary compute. The load on the application is irregular. The application experiences long periods of no usage, followed by sudden and significant increases and decreases in traffic. The application is write-heavy and stores data in an Amazon Aurora MySQL database. The database runs on an Amazon RDS memory optimized DB instance that is not able to handle the load.

What is the MOST cost-effective way for the company to handle the sudden and significant changes in traffic?

- A. Add additional read replicas to the database. Purchase Instance Savings Plans and RDS Reserved Instances.
- B. Migrate the database to an Aurora multi-master DB cluster. Purchase Instance Savings Plans.
- C. Migrate the database to an Aurora global database. Purchase Compute Savings Plans and RDS Reserved Instances.
- D. Migrate the database to Aurora Serverless v1. Purchase Compute Savings Plans.
- E. Migrate the database to Aurora Serverless v2. Purchase Compute Savings Plans only for predictable baseline Lambda and AWS Fargate usage.

ANSWER: E

Explanation:

Migrate the database to Aurora Serverless v2 is the best fit because the workload is highly variable, write-heavy, and has long idle periods followed by rapid traffic spikes. Aurora Serverless v2 automatically adjusts database capacity in fine-grained increments based on application demand, which helps the database scale for sudden increases in write activity without permanently provisioning a large memory-optimized instance. For compatible Aurora MySQL versions, Aurora Serverless v2 can also support automatic pause and resume with a minimum capacity of 0 ACUs, which is especially cost-effective when the application has long periods with no usage. Because the application compute layer uses AWS Lambda and AWS Fargate, Compute Savings Plans can be considered only for any predictable baseline usage; they apply to both Lambda and Fargate, but commitments should match expected steady consumption to remain cost-effective. See the AWS documentation for [Aurora Serverless v2](#) and [AWS Savings Plans](#).

QUESTION NO: 2

A solutions architect has implemented a SAML 2.0 federated identity solution with their company's on-premises identity provider (IdP) to authenticate users' access to the AWS environment. When the solutions architect tests authentication through the federated identity web portal, access to the AWS environment is granted. However, when test users attempt to authenticate through the federated identity web portal, they are not able to access the AWS environment. Which items should the solutions architect check to ensure identity federation is properly configured? (Choose three.)

- A. The IAM user's permissions policy has allowed the use of SAML federation for that user.
- B. The IAM roles created for the federated users' or federated groups' trust policy have set the SAML provider as the principal.
Test users are not in the AWSFederatedUsers group in the company's IdP.
- C. The web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider, the ARN of the IAM role, and the SAML assertion from IdP.
- D. The on-premises IdP's DNS hostname is reachable from the AWS environment VPCs.
- E. The company's IdP defines SAML assertions that properly map users or groups. In the company to IAM roles with appropriate permissions.

ANSWER: B C E

Explanation:

For SAML 2.0 federation to AWS, the IAM role that federated users assume must trust the IAM SAML identity provider. That trust relationship allows AWS STS to accept a valid SAML assertion from the configured provider and issue temporary credentials for the role. The application or federation portal must also submit the SAML assertion to AWS STS by using AssumeRoleWithSAML, including the SAML provider ARN and the target IAM role ARN. This is the core exchange that converts a successful IdP authentication into AWS temporary security credentials. Finally, the IdP must emit the correct SAML attributes that map users or groups to the intended IAM roles, and users must belong to the groups or match the rules that cause those role attributes to be included. If the architect can authenticate but test users cannot, validating these role mappings and group-based assertions is essential. AWS documents this flow in its guidance for [SAML 2.0 federation](#) and in the [AssumeRoleWithSAML API reference](#).

QUESTION NO: 3

A company uses AWS Organizations to manage more than 1.000 AWS accounts. The company has created a new developer organization. There are 540 developer member accounts that must be moved to the new developer organization. All accounts are set up with all the required information so that each account can be operated as a standalone account.

Which combination of steps should a solutions architect take to move all of the developer accounts to the new developer organization? (Select THREE)

- A. Call the MoveAccount operation in the Organizations API from the old organization's management account to migrate the developer accounts to the new developer organization
- B. From the management account remove each developer account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API
- C. From each developer account, remove the account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API
- D. Sign in to the new developer organization's management account and create a placeholder member account that acts as a target for the developer account migration
- E. Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization's management account to send invitations to the developer accounts.
- F. Have each developer sign in to their account and confirm to join the new developer organization.

ANSWER: B E F

Explanation:

The correct migration flow is to first remove each developer account from the existing organization, then invite each standalone account into the new developer organization, and finally have each invited account accept the invitation. "From the management account remove each developer account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API" is correct because AWS Organizations requires the old organization's management account to remove member accounts, and the accounts must have the required standalone account information before removal. "Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization's management account to send invitations to the developer accounts." is correct because the new organization's management account initiates the join process by creating invitations, which are represented as handshakes. "Have each developer sign in to their account and confirm to join the new developer organization." is correct because each invited account must accept the handshake before becoming a member of the new organization. This sequence aligns with the AWS Organizations account lifecycle for moving accounts between separate organizations. See the AWS API references for [RemoveAccountFromOrganization](#) and [InviteAccountToOrganization](#).

QUESTION NO: 4

A company is migrating its three-tier web application from on-premises to the AWS Cloud. The company has the following requirements for the migration process:

- Ingest machine images from the on-premises environment.
- Synchronize changes from the on-premises environment to the AWS environment until the production cutover.

- Minimize downtime when executing the production cutover.
- Migrate the virtual machines' root volumes and data volumes.

Which solution will satisfy these requirements with minimal operational overhead?

- A.** Use AWS Server Migration Service (SMS) to create and launch a replication job for each tier of the application. Launch instances from the AMIs created by AWS SMS. After initial testing, perform a final replication and create new instances from the updated AMIs.
- B.** Create an AWS CLIVM Import/Export script to migrate each virtual machine. Schedule the script to run incrementally to maintain changes in the application. Launch instances from the AMIs created by VM Import/Export. Once testing is done, rerun the script to do a final import and launch the instances from the AMIs.
- C.** Use AWS Server Migration Service (SMS) to upload the operating system volumes. Use the AWS CLI import-snaps hot command for the data volumes. Launch instances from the AMIs created by AWS SMS and attach the data volumes to the instances. After initial testing, perform a final replication, launch new instances from the replicated AMIs, and attach the data volumes to the instances.
- D.** Use AWS Application Discovery Service and AWS Migration Hub to group the virtual machines as an application. Use the AWS CLI VM Import/Export script to import the virtual machines as AMIs. Schedule the script to run incrementally to maintain changes in the application. Launch instances from the AMIs. After initial testing, perform a final virtual machine import and launch new instances from the AMIs.
- E.** Use AWS Application Migration Service to install the replication agent on each source virtual machine. Configure launch settings, perform non-disruptive tests, and execute a cutover to launch migrated instances with the replicated root and data volumes.

ANSWER: E

Explanation:

Use AWS Application Migration Service to install the replication agent on each source virtual machine. Configure launch settings, perform non-disruptive tests, and execute a cutover to launch migrated instances with the replicated root and data volumes. is correct because AWS Application Migration Service (AWS MGN) is the current AWS service designed for large-scale lift-and-shift migrations with minimal downtime and low operational effort. After the AWS Replication Agent is installed on each source server, AWS MGN continuously replicates block-level data from the source machines to a low-cost staging area in AWS. This ongoing replication keeps the AWS copy synchronized until the cutover window, including the operating system disk and selected additional data disks. The service also lets teams perform non-disruptive test launches before production cutover, then initiate a cutover action that launches the migrated EC2 instances from the latest replicated state. This workflow directly supports testing, final synchronization, and a short cutover window without requiring custom import scripts or separate handling of root and data volumes. AWS describes Application Migration Service as the recommended service for lift-and-shift migrations to AWS. See [AWS Application Migration Service documentation](#) and the [AWS Application Migration Service product page](#).

QUESTION NO: 5

A company creates an AWS Control Tower landing zone to manage and govern a multi-account AWS environment. The company's security team will deploy preventive controls and detective controls to monitor AWS services across all the accounts. The security team needs a centralized view of the security state of all the accounts.

Which solution will meet these requirements?

- A.** From the AWS Control Tower management account, use AWS CloudFormation StackSets to deploy an AWS Config conformance pack to all accounts in the organization
- B.** Enable Amazon Detective for the organization in AWS Organizations Designate one AWS account as the delegated administrator for Detective
- C.** From the AWS Control Tower management account, deploy an AWS CloudFormation stack set that uses the automatic deployment option to enable Amazon Detective for the organization
- D.** Enable AWS Security Hub for the organization in AWS Organizations Designate one AWS account as the delegated administrator for Security Hub

ANSWER: D

Explanation:

Enable AWS Security Hub for the organization in AWS Organizations Designate one AWS account as the delegated administrator for Security Hub

is correct because AWS Security Hub is designed to provide a centralized, organization-wide view of security findings and security posture across multiple AWS accounts and Regions. When integrated with AWS Organizations, Security Hub can be centrally enabled and managed, and a delegated administrator account can administer member accounts without requiring day-to-day security operations to run from the management account. Security Hub aggregates findings from AWS services such as Amazon GuardDuty, Amazon Inspector, IAM Access Analyzer, AWS Config, and supported partner products, and it evaluates accounts against security standards and controls. This aligns directly with the requirement for the security team to monitor preventive and detective controls across all accounts and have a centralized view of security state. In an AWS Control Tower environment, using Security Hub with AWS Organizations is a common best-practice pattern for centralized security governance and visibility. See the AWS documentation for [managing Security Hub accounts with AWS Organizations](#) and the AWS Control Tower guidance for [Security Hub integration](#).

QUESTION NO: 6

A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand. Which solutions meet these requirements? (Choose two.)

- A.** Create an Amazon API Gateway REST API. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.
- B.** Create an Amazon API Gateway HTTP API. Configure this API with direct integrations to Dynamo DB by using API Gateway's AWS integration type.
- C.** Create an Amazon API Gateway HTTP API. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.
- D.** Create an accelerator in AWS Global Accelerator. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.
- E.** Create a Network Load Balancer. Configure listener rules to forward requests to the appropriate AWS Lambda functions.

ANSWER: A C

Explanation:

Create an Amazon API Gateway REST API. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type is correct because Amazon API Gateway REST APIs can use AWS service integrations to call DynamoDB actions directly over HTTPS without managing servers. API Gateway provides the public HTTPS endpoint, handles request routing and throttling, and scales automatically, while DynamoDB also provides fully managed, highly scalable data storage. This is a good fit when the API can map requests directly to DynamoDB operations by using API Gateway mapping templates and IAM-based service integration permissions. See the AWS documentation for [API Gateway REST API integration types](#).

Create an Amazon API Gateway HTTP API. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables is also correct. HTTP APIs provide a lightweight, publicly accessible HTTPS API endpoint, and AWS Lambda functions can implement the application logic to query or scan DynamoDB tables and return responses. Lambda and DynamoDB both scale automatically based on demand, so the overall architecture remains serverless and elastic. This pattern is commonly used when the API requires custom logic, validation, response formatting, or access to multiple DynamoDB tables. See AWS guidance on [HTTP API Lambda integrations](#).

QUESTION NO: 7

A large company recently experienced an unexpected increase in Amazon RDS and Amazon DynamoDB costs The company needs to increase visibility into details of AWS Billing and Cost Management There are various accounts

associated with AWS Organizations, including many development and production accounts. There is no consistent tagging strategy across the organization, but there are guidelines in place that require all infrastructure to be deployed using AWS Cloud Formation with consistent tagging Management requires cost center numbers and project ID numbers for all existing and future DynamoDB tables and RDS instances

Which strategy should the solutions architect provide to meet these requirements?

- A.** Use Tag Editor to tag existing resources Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources
- B.** Use an AWS Config rule to alert the finance team of untagged resources Create a centralized AWS Lambda based solution to tag untagged RDS databases and DynamoDB resources every hour using a cross-account role.
- C.** Use Tag Editor to tag existing resources Create cost allocation tags to define the cost center and project ID Use SCPs to restrict resource creation that do not have the cost center and project ID on the resource.
- D.** Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources Update existing federated roles to restrict privileges to provision resources that do not include the cost center and project ID on the resource

ANSWER: C

Explanation:

Use Tag Editor to tag existing resources Create cost allocation tags to define the cost center and project ID Use SCPs to restrict resource creation that do not have the cost center and project ID on the resource. is correct because it addresses both the historical cleanup requirement and the future governance requirement across AWS Organizations. AWS Tag Editor can be used to locate and apply required tags, such as cost center and project ID, to existing supported resources so that current RDS instances and DynamoDB tables can be brought into compliance. After the tag keys are activated as user-defined cost allocation tags in AWS Billing and Cost Management, AWS can use those tags to categorize and report costs in billing tools such as Cost Explorer and cost allocation reports. For future resources, service control policies provide an organization-wide preventive guardrail. An SCP can deny resource creation API calls when required request tags are missing, which is the right control when many development and production accounts must follow the same tagging rules. This approach fits the stated CloudFormation-based deployment model because templates can include the required tags, while noncompliant creates are blocked centrally. See AWS guidance on [cost allocation tags](#) and [SCP examples for tagging requirements](#).

QUESTION NO: 8

A company has developed a new release of a popular video game and wants to make it available for public download. The new release package is approximately 5 GB in size. The company provides downloads for existing releases from a Linux-based, publicly facing FTP site hosted in an on-premises data center. The company expects the new release will be downloaded by users worldwide. The company wants a solution that provides improved download performance and low transfer costs, regardless of a user's location.

- A.** Store the game files on Amazon EBS volumes mounted on Amazon EC2 instances within an Auto Scaling group. Configure an FTP service on the EC2 instances. Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to download the package.
- B.** Store the game files on Amazon EFS volumes that are attached to Amazon EC2 instances within an Auto Scaling group. Configure an FTP service on each of the EC2 instances. Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to download the package.
- C.** Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Use Amazon CloudFront for the website. Publish the game download URL for users to download the package.
- D.** Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Set Requester Pays for the S3 bucket. Publish the game download URL for users to download the package.

ANSWER: C

Explanation:

Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Use Amazon CloudFront for the website. Publish the game download URL for users to download the package is correct because it uses managed AWS services that are purpose-built for global, high-scale content distribution. Amazon S3 provides durable object storage for the 5 GB game package and can serve the file over HTTP/HTTPS without the company managing FTP servers, storage replication, or instance capacity. Amazon CloudFront caches content at edge locations around the world, so users download the package from a nearby edge location where possible, reducing latency and improving throughput for a global audience. CloudFront also helps reduce origin load and can lower delivery costs compared with repeatedly serving large files from a single origin or from self-managed compute infrastructure. Amazon Route 53 can provide DNS for the public download hostname and route users to the CloudFront distribution. This architecture is highly available, scalable, and operationally simpler for public software distribution. See the AWS guidance on [Amazon CloudFront](#) and using [Amazon S3 static website hosting](#).

QUESTION NO: 9

A company has a data lake in Amazon S3 that needs to be accessed by hundreds of applications across many AWS accounts. The company's information security policy states that the S3 bucket must not be accessed over the public internet and that each application should have the minimum permissions necessary to function.

To meet these requirements, a solutions architect plans to use an S3 access point that is restricted to specific VPCs for each application.

Which combination of steps should the solutions architect take to implement this solution? (Select TWO.)

- A.** Create an S3 access point for each application in the AWS account that owns the S3 bucket. Configure each access point to be accessible only from the application's VPC. Update the bucket policy to require access from an access point.
- B.** Create an interface endpoint for Amazon S3 in each application's VPC. Configure the endpoint policy to allow access to an S3 access point. Create a VPC gateway attachment for the S3 endpoint.
- C.** Create a gateway endpoint for Amazon S3 in each application's VPC. Configure the endpoint policy to allow access to an S3 access point. Specify the route table that is used to access the access point.
- D.** Create an S3 access point for each application in each AWS account and attach the access points to the S3 bucket. Configure each access point to be accessible only from the application's VPC. Update the bucket policy to require access from an access point.
- E.** Create a gateway endpoint for Amazon S3 in the data lake's VPC. Attach an endpoint policy to allow access to the S3 bucket. Specify the route table that is used to access the bucket.

ANSWER: C D

Explanation:

The correct implementation is to create an S3 access point for each application in each AWS account and attach the access points to the S3 bucket, together with a gateway endpoint for Amazon S3 in each application's VPC. A VPC-restricted S3 access point limits access so requests must originate from the configured VPC, and separate access points allow each application to receive its own narrowly scoped permissions through access point policies and the central bucket policy. This is a good fit for a shared data lake because the bucket owner can delegate access through access points while still requiring all access to use approved access point paths.

Creating a gateway endpoint for Amazon S3 in each application's VPC keeps traffic on the AWS network instead of traversing the public internet. The route tables associated with the application subnets must point S3 traffic to the gateway endpoint, and the endpoint policy can further restrict access to the intended S3 access point resources. AWS documents S3 access points and VPC-restricted access patterns in the [Amazon S3 access points guide](#) and describes private S3 connectivity with gateway endpoints in the [Amazon S3 gateway endpoint documentation](#).

QUESTION NO: 10

A weather service provides high-resolution weather maps from a web application hosted on AWS in the eu-west-1 Region. The weather maps are updated frequently and stored in Amazon S3 along with static HTML content. The web application is fronted by Amazon CloudFront.

The company recently expanded to serve users in the us-east-1 Region, and these new users report that viewing their respective weather maps is slow from time to time.

Which combination of steps will resolve the us-east-1 performance issues? (Choose two.)

- A. Configure the AWS Global Accelerator endpoint for the S3 bucket in eu-west-1. Configure endpoint groups for TCP ports 80 and 443 in us-east-1.
- B. Create a new S3 bucket in us-east-1. Configure S3 cross-Region replication to synchronize from the S3 bucket in eu-west-1. Most Voted
- C. Use Lambda@Edge to modify requests from North America to use the S3 Transfer Acceleration endpoint in us-east-1.
- D. Use Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1. Most Voted
- E. Configure the AWS Global Accelerator endpoint for us-east-1 as an origin on the CloudFront distribution. Use Lambda@Edge to modify requests from North America to use the new origin.

ANSWER: B D

Explanation:

The correct combination is to create a new S3 bucket in us-east-1 and configure S3 cross-Region replication to synchronize from the S3 bucket in eu-west-1, then use Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1. Because the weather maps are updated frequently, CloudFront edge caches may experience periodic cache misses or revalidations that require fetching content from the origin. If the only origin is the S3 bucket in eu-west-1, North American viewers can see intermittent latency when CloudFront must retrieve fresh objects from Europe. Replicating the objects into an S3 bucket in us-east-1 places the origin closer to those users and reduces origin fetch latency.

Lambda@Edge can dynamically inspect viewer request attributes, such as CloudFront viewer location headers, and rewrite the request origin so North American requests are served from the us-east-1 S3 origin while other users can continue to use the existing eu-west-1 origin. This pattern combines regional origin placement with CloudFront's global edge network for better performance during cache misses and frequent updates. See AWS documentation for [Amazon S3 replication](#) and [Lambda@Edge](#).

QUESTION NO: 11

A company is running a three-tier web application in an on-premises data center. The frontend is served by an Apache web server, the middle tier is a monolithic Java application, and the storage tier is a PostgreSQL database.

During a recent marketing promotion, customers could not place orders through the application because the application crashed. An analysis showed that all three tiers were overloaded. The application became unresponsive, and the database reached its capacity limit because of read operations. The company already has several similar promotions scheduled in the near future.

A solutions architect must develop a plan for migration to AWS to resolve these issues. The solution must maximize scalability and must minimize operational effort.

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

- A. Refactor the frontend so that static assets can be hosted on Amazon S3. Use Amazon CloudFront to serve the frontend to customers. Connect the frontend to the Java application.
- B. Rehost the Apache web server of the frontend on Amazon EC2 instances that are in an Auto Scaling group. Use a load balancer in front of the Auto Scaling group. Use Amazon Elastic File System (Amazon EFS) to host the static assets that the Apache web server needs.
- C. Rehost the Java application in an AWS Elastic Beanstalk environment that includes auto scaling.
- D. Refactor the Java application. Develop a Docker container to run the Java application. Use AWS Fargate to host the container.
- E. Use AWS Database Migration Service (AWS DMS) to replatform the PostgreSQL database to an Amazon Aurora PostgreSQL database. Use Aurora Auto Scaling for read replicas.

F. Rehost the PostgreSQL database on an Amazon EC2 instance that has twice as much memory as the on-premises server.

ANSWER: A C E

Explanation:

Refactor the frontend so that static assets can be hosted on Amazon S3. Use Amazon CloudFront to serve the frontend to customers. Connect the frontend to the Java application is correct because S3 and CloudFront provide highly scalable, managed delivery of static content with minimal operational overhead, reducing load on the application tier during traffic spikes. Rehost the Java application in an AWS Elastic Beanstalk environment that includes auto scaling is correct because Elastic Beanstalk supports Java applications and manages capacity provisioning, load balancing, scaling, deployment, and health monitoring while still allowing the monolithic application to run with limited refactoring. Use AWS Database Migration Service (AWS DMS) to replatform the PostgreSQL database to an Amazon Aurora PostgreSQL database. Use Aurora Auto Scaling for read replicas is correct because Aurora PostgreSQL is a managed relational database service that is compatible with PostgreSQL, and Aurora Auto Scaling can add or remove Aurora Replicas to handle changes in read demand. This directly addresses the database read-capacity bottleneck while minimizing database administration effort. See [Amazon CloudFront with Amazon S3 origins](#) and [Aurora Auto Scaling with Aurora Replicas](#).

QUESTION NO: 12

A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.

A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.

Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

- A.** Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.
- B.** Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enabled. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template. Use the EBS volume as the shared storage for the duration of the job. Detach the EBS volume when the job is complete.
- C.** Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.
- D.** Migrate the data from the existing shared file system to an Amazon S3 bucket. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the job. Delete the file gateway when the job is complete.

ANSWER: A

Explanation:

Migrating the data to Amazon S3 with S3 Intelligent-Tiering and creating a temporary Amazon FSx for Lustre file system by using lazy loading is the best fit. S3 provides low-cost durable storage for the 200 TB dataset when the monthly job is not running, and S3 Intelligent-Tiering can reduce ongoing storage costs for objects with changing or unpredictable access patterns. FSx for Lustre is designed for high-performance workloads that need shared, POSIX-compatible file access from large EC2 fleets. By creating the FSx for Lustre file system only for the 72-hour monthly processing window, the company

pays for the high-performance file system only while it is needed. Lazy loading is especially cost-effective here because the job reads only a subset of the files; FSx can present the S3 namespace immediately and load file data from S3 as files are accessed, avoiding the need to preload all 200 TB before each run. This combination provides durable low-cost storage between runs and high-throughput shared file access during the run. See [Amazon FSx for Lustre documentation](#) and [Amazon S3 Intelligent-Tiering](#).

QUESTION NO: 13

A company stores application data in many Amazon S3 buckets in one AWS account. Some of the S3 buckets contain sensitive data. The company does not have data inventory for the S3 buckets. The company uses server-side encryption with Amazon S3 managed keys (SSE-S3) to encrypt all data in the S3 buckets.

A solutions architect must design a solution to encrypt sensitive data with a key that only administrators can access.

Which solution will meet these requirements?

- A.** Use Amazon Inspector to determine which S3 buckets contain sensitive data. Create a new AWS KMS customer managed key and a key policy that provides access to administrators only. Set default S3 bucket encryption to use the new KMS key (SSE-KMS). Update the S3 bucket policy to add a Deny effect and a Condition element of "StringNotEquals" : { "s3:x-amz-server-side-encryption" : "aws:kms" }.
- B.** Use Amazon Inspector to determine which S3 buckets contain sensitive data. Update the key policy on the AWS managed key to provide access to administrators only. Use AWS Batch to encrypt all existing objects that include sensitive data in the S3 buckets with the updated AWS managed key.
- C.** Use Amazon Macie to determine which S3 buckets contain sensitive data. Create a new AWS KMS customer managed key and a key policy that provides access to administrators only. Set default S3 bucket encryption to use the new KMS key (SSE-KMS). Create an AWS Step Functions workflow to encrypt all existing S3 objects that include sensitive data by using the new KMS key.
- D.** Use Amazon Macie to determine which S3 buckets contain sensitive data. Update the key policy on the AWS managed key to provide access to administrators only. Update the S3 bucket policy to add a Deny effect and a Condition element of "StringNotEquals" : { "s3:x-amz-server-side-encryption" : "aws:kms" }.

ANSWER: C

Explanation:

The correct solution is the one that uses Amazon Macie to identify sensitive data in Amazon S3, creates a new AWS KMS customer managed key with a restrictive key policy for administrators, configures the relevant S3 buckets to use SSE-KMS by default, and re-encrypts existing sensitive objects with the new key. Amazon Macie is designed to discover and classify sensitive data stored in S3, which directly addresses the lack of an existing data inventory. A customer managed AWS KMS key is appropriate because its key policy can be controlled by the company, allowing the organization to restrict key administration and key usage to specific administrator principals. Setting default bucket encryption to SSE-KMS ensures that newly written objects use the new KMS key. Because changing default encryption does not automatically re-encrypt existing objects, a workflow to copy or rewrite the existing sensitive objects with the new KMS key is also required. AWS Step Functions can orchestrate that re-encryption process across the discovered objects. See [Amazon Macie documentation](#) and [Amazon S3 SSE-KMS documentation](#).

QUESTION NO: 14

A company runs many workloads on AWS and uses AWS Organizations to manage its accounts. The workloads are hosted on Amazon EC2, AWS Fargate, and AWS Lambda. Some of the workloads have unpredictable demand. Accounts record high usage in some months and low usage in other months.

The company wants to optimize its compute costs over the next 3 years. A solutions architect obtains a 6-month average for each of the accounts across the organization to calculate usage.

Which solution will provide the MOST cost savings for all the organization's compute usage?

- A.** Purchase Reserved Instances for the organization to match the size and number of the most common EC2 instances from the member accounts.
- B.** Purchase a Compute Savings Plan for the organization from the management account by using the recommendation at the management account level.
- C.** Purchase Reserved Instances for each member account that had high EC2 usage according to the data from the last 6 months.
- D.** Purchase an EC2 Instance Savings Plan for each member account from the management account based on EC2 usage data from the last 6 months.

ANSWER: B

Explanation:

Purchase a Compute Savings Plan for the organization from the management account by using the recommendation at the management account level is correct because Compute Savings Plans provide the broadest compute discount coverage across Amazon EC2, AWS Fargate, and AWS Lambda. This matches the company's mixed workload environment and is especially useful when demand varies by account or service over time. In AWS Organizations, the management account can view recommendations that are calculated across the organization's consolidated usage, which helps size the commitment based on aggregate historical usage rather than isolated account-level peaks. Savings Plans also automatically apply to eligible usage across linked accounts when sharing is enabled, helping the company maximize utilization of the hourly commitment over the 3-year term. For a business with unpredictable monthly usage patterns, an organization-level Compute Savings Plan based on recommendations is typically more flexible and cost-effective than committing to specific EC2 instance attributes. AWS documentation states that Compute Savings Plans apply to EC2 usage regardless of instance family, size, Availability Zone, Region, operating system, or tenancy, and also apply to Fargate and Lambda usage. See [AWS Savings Plans](#) and [How Savings Plans apply to usage](#).

QUESTION NO: 15

A company has applications in an AWS account that is named Source. The account is in an organization in AWS Organizations. One of the applications uses AWS Lambda functions and stores inventory data in an Amazon Aurora database. The application deploys the Lambda functions by using a deployment package. The company has configured automated backups for Aurora. The company wants to migrate the Lambda functions and the Aurora database to a new AWS account that is named Target. The application processes critical data, so the company must minimize downtime. Which solution will meet these requirements?

- A.** Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the automated Aurora DB cluster snapshot with the Target account.
- B.** Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the Aurora DB cluster with the Target account by using AWS Resource Access Manager (AWS RAM). Grant the Target account permission to clone the Aurora DB cluster.
- C.** Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions and the Aurora DB cluster with the Target account. Grant the Target account permission to clone the Aurora DB cluster.
- D.** Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions with the Target account. Share the automated Aurora DB cluster snapshot with the Target account.

ANSWER: B

Explanation:

The correct solution is to download the Lambda function deployment package from the Source account, use it to create new Lambda functions in the Target account, and share the Aurora DB cluster with the Target account by using AWS Resource Access Manager so the Target account can clone the Aurora DB cluster. Lambda functions themselves are not migrated across accounts by AWS RAM as shared resources, so recreating the functions from the deployment package is the appropriate account-migration approach for the compute layer. For the database layer, Aurora supports cross-account cloning through AWS RAM. A cloned Aurora DB cluster uses a copy-on-write mechanism, which is much faster than restoring a full snapshot and is well suited when downtime must be minimized for a critical application. The Target account

can create its own Aurora clone from the shared cluster and then the application can be cut over after validation. This aligns with AWS guidance for Aurora cloning and supported RAM sharing workflows. References: [Amazon Aurora cloning](#) and [AWS Resource Access Manager](#).

QUESTION NO: 16

A company is hosting a critical application on a single Amazon EC2 instance. The application uses an Amazon ElastiCache for Redis single-node cluster for an in-memory data store. The application uses an Amazon RDS for MariaDB DB instance for a relational database. For the application to function, each piece of the infrastructure must be healthy and must be in an active state.

A solutions architect needs to improve the application's architecture so that the infrastructure can automatically recover from failure with the least possible downtime.

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

A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.

Using an Elastic Load Balancer (ELB) to distribute traffic across multiple EC2 instances and ensuring that the EC2 instances are part of an Auto Scaling group with a minimum capacity of two instances will improve the availability and scalability of the application. The ELB will automatically route traffic to healthy instances, and the Auto Scaling group will automatically scale the number of instances based on demand. In the event of a failure, the Auto Scaling group will automatically replace the failed instance, minimizing downtime.

B. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are configured in unlimited mode.

C. Modify the DB instance to create a read replica in the same Availability Zone. Promote the read replica to be the primary DB instance in failure scenarios.

D. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.

As mentioned in the previous answer, a Multi-AZ deployment for Amazon RDS for MariaDB DB instances provides enhanced availability and failover support for DB instances. In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous "standby" replica in a different availability zone (AZ) than the primary DB instance. In the event of planned or unplanned outages, Amazon RDS performs an automatic failover to the standby, minimizing downtime.

E. Create a replication group for the ElastiCache for Redis cluster. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.

Creating a replication group for the ElastiCache for Redis cluster and configuring the cluster to use an Auto Scaling group that has a minimum capacity of two instances will improve the availability and scalability of the ElastiCache service. The replication group will provide failover support for the Redis cluster and the Auto Scaling group will automatically increase capacity in the event of a failure, minimizing downtime.

Reference:

F. Create a replication group for the ElastiCache for Redis cluster. Enable Multi-AZ on the cluster.

ANSWER: A D F

Explanation:

The correct combination is to use an Elastic Load Balancer with multiple EC2 instances in an Auto Scaling group with a minimum capacity of two instances, modify the RDS for MariaDB DB instance to use a Multi-AZ deployment, and create a replication group for ElastiCache for Redis with Multi-AZ enabled. Together, these changes remove the single points of failure across the application tier, relational database tier, and in-memory cache tier. Elastic Load Balancing routes requests only to healthy EC2 instances, while Amazon EC2 Auto Scaling can maintain the required number of running instances and replace unhealthy instances automatically. Amazon RDS Multi-AZ provides a synchronously replicated standby in another Availability Zone and performs automatic failover if the primary DB instance becomes unavailable. For Redis, an ElastiCache replication group with Multi-AZ and automatic failover provides a primary node and replicas across Availability Zones, allowing ElastiCache to promote a replica if the primary node fails. This architecture provides automated recovery with minimal downtime across all required components. References: [Amazon EC2 Auto Scaling benefits](#) and [ElastiCache for Redis Multi-AZ with automatic failover](#).

QUESTION NO: 17

A video streaming company recently launched a mobile app for video sharing. The app uploads various files to an Amazon S3 bucket in the us-east-1 Region. The files range in size from 1 GB to 10 GB. Users who access the app from Australia have experienced uploads that take long periods of time. Sometimes the files fail to completely upload for these users. A solutions architect must improve the app's performance for these uploads. Which solutions will meet these requirements? (Choose two.)

- A. Enable S3 Transfer Acceleration on the S3 bucket. Configure the app to use the Transfer Acceleration endpoint for uploads.
- B. Configure an S3 bucket in each Region to receive the uploads. Use S3 Cross-Region Replication to copy the files to the distribution S3 bucket.
- C. Set up Amazon Route 53 with latency-based routing to route the uploads to the nearest S3 bucket Region.
- D. Configure the app to break the video files into chunks. Use a multipart upload to transfer files to Amazon S3.
- E. Modify the app to add random prefixes to the files before uploading.

ANSWER: A D

Explanation:

Enable S3 Transfer Acceleration on the S3 bucket. Configure the app to use the Transfer Acceleration endpoint for uploads. is correct because S3 Transfer Acceleration is designed to speed up long-distance transfers into S3 by routing traffic through AWS edge locations and the AWS global network. For users in Australia uploading to a bucket in us-east-1, this can reduce latency and improve throughput without requiring the application to manage regional buckets. AWS specifically recommends this feature for transferring data over long geographic distances to a centralized S3 bucket. See the AWS documentation for [Amazon S3 Transfer Acceleration](#).

Configure the app to break the video files into chunks. Use a multipart upload to transfer files to Amazon S3. is also correct. Multipart upload is a best practice for large objects, especially for files in the 1 GB to 10 GB range. It allows parts to be uploaded independently and in parallel, which can improve throughput. It also improves resilience because a failed part can be retried without restarting the entire upload. AWS recommends multipart upload for large objects and unreliable networks. See [Uploading and copying objects using multipart upload](#).

QUESTION NO: 18

A financial company is planning to migrate its web application from on premises to AWS. The company uses a third-party security tool to monitor the inbound traffic to the application. The company has used the security tool for the last 15 years, and the tool has no cloud solutions available from its vendor. The company's security team is concerned about how to integrate the security tool with AWS technology. The company plans to deploy the application migration to AWS on Amazon EC2 instances. The EC2 instances will run in an Auto Scaling group in a dedicated VPC. The company needs to use the security tool to inspect all packets that come in and out of the VPC. This inspection must occur in real time and must not affect the application's performance. A solutions architect must design a target architecture on AWS that is highly available within an AWS Region. Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

- A. Deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC
- B. Deploy the web application behind a Network Load Balancer
- C. Deploy an Application Load Balancer in front of the security tool instances
- D. Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool
- E. Provision a transit gateway to facilitate communication between VPCs.

ANSWER: A D

Explanation:

The correct combination is to deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC and provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool. Gateway Load Balancer is designed specifically for transparent, inline deployment of third-party virtual appliances such as firewalls, intrusion detection and prevention systems, and packet inspection tools. It operates at layer 3, uses the GENEVE protocol to encapsulate traffic, and preserves the original packet flow so appliances can inspect traffic without requiring the application to be aware of the inspection layer.

Running the security tool on EC2 instances in an Auto Scaling group provides elasticity and high availability for the inspection fleet. Combining this with Gateway Load Balancer across Availability Zones gives a highly available regional architecture and distributes traffic to healthy appliance instances while allowing real-time packet inspection at scale. This pattern is AWS's recommended architecture for inserting security appliances transparently into VPC traffic flows. See the AWS documentation for [Gateway Load Balancers](#) and the [AWS networking inspection deployment patterns](#).

QUESTION NO: 19

A company is running a traditional web application on Amazon EC2 instances. The company needs to refactor the application as microservices that run on containers. Separate versions of the application exist in two distinct environments: production and testing. Load for the application is variable, but the minimum load and the maximum load are known. A solutions architect needs to design the updated application with a serverless architecture that minimizes operational complexity.

Which solution will meet these requirements MOST cost-effectively?

- A.** Upload the container images to AWS Lambda as functions. Configure a concurrency limit for the associated Lambda functions to handle the expected peak load. Configure two separate Lambda integrations within Amazon API Gateway: one for production and one for testing.
- B.** Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Container Service (Amazon ECS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the ECS clusters.
- C.** Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Kubernetes Service (Amazon EKS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the EKS clusters.
- D.** Upload the container images to AWS Elastic Beanstalk. In Elastic Beanstalk, create separate environments and deployments for production and testing. Configure two separate Application Load Balancers to direct traffic to the Elastic Beanstalk deployments.

ANSWER: B

Explanation:

Uploading the container images to Amazon Elastic Container Registry and running them on Amazon Elastic Container Service with the Fargate launch type is the best fit for a serverless container architecture. Amazon ECS with AWS Fargate removes the need to provision, patch, scale, or manage EC2 container hosts, while still allowing the application to run as containerized microservices. Because the minimum and maximum load are known, ECS Service Auto Scaling can be configured to scale the number of running tasks between defined bounds, helping match capacity to demand and control cost. Using separate ECS deployments for production and testing provides environment isolation, and Application Load Balancers are a standard way to route HTTP or HTTPS traffic to ECS services running on Fargate. This design keeps operational overhead low while paying for the vCPU and memory resources consumed by running tasks rather than managing a fixed fleet of servers. AWS documents Fargate as a serverless compute engine for containers with ECS, and ECS integrates directly with Elastic Load Balancing and service auto scaling. References: [AWS Fargate for Amazon ECS](#) and [Amazon ECS service auto scaling](#).

QUESTION NO: 20

A company is updating an application that customers use to make online orders. The number of attacks on the application by bad actors has increased recently. The company will host the updated application on an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use Amazon DynamoDB to store application data. A public Application Load Balancer (ALB) will provide end users with access to the application. The company must prevent attacks and ensure

business continuity with minimal service interruptions during an ongoing attack. Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

- A. Create an Amazon CloudFront distribution with the ALB as the origin. Add a custom header and random value on the CloudFront domain. Configure the ALB to conditionally forward traffic if the header and value match.
- B. Deploy the application in two AWS Regions. Configure Amazon Route 53 to route to both Regions with equal weight.
- C. Configure auto scaling for Amazon ECS tasks Create a DynamoDB Accelerator (DAX) cluster.
- D. Configure Amazon ElastiCache to reduce overhead on DynamoDB.
- E. Deploy an AWS WAF web ACL that includes an appropriate rule group. Associate the web ACL with the Amazon CloudFront distribution.

ANSWER: A E

Explanation:

Creating an Amazon CloudFront distribution with the ALB as the origin and using a custom header with a random value is correct because it helps ensure that legitimate application traffic reaches the ALB through CloudFront rather than directly from the internet. The ALB can use listener rules to forward only requests that contain the expected header value, which helps prevent attackers from bypassing CloudFront-level protections. CloudFront also improves availability during attacks by absorbing and distributing traffic across AWS edge locations. AWS documents this origin-protection pattern for custom headers in CloudFront: [Add custom headers to origin requests](#).

Deploying an AWS WAF web ACL with an appropriate rule group and associating it with the CloudFront distribution is also correct. AWS WAF can inspect HTTP and HTTPS requests before they reach the application, using managed rule groups, rate-based rules, IP reputation lists, and application-specific rules to block common web exploits and abusive request patterns. Associating AWS WAF with CloudFront provides centralized edge filtering, reducing malicious traffic before it reaches the public ALB and ECS services. This combination is cost-effective because it uses managed edge protections without requiring multi-Region active-active deployment or additional caching layers solely for attack mitigation. See AWS guidance here: [Using AWS WAF with CloudFront](#).

QUESTION NO: 21

A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.

What should a solutions architect do to meet these requirements?

- A. Create a new developer account. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organizations. Enforce a tagging policy that denotes Region affinity.
- B. Create an SCP that denies the launch of all EC2 instances except I3.small EC2 instances in us-east-2. Attach the SCP to the project's account.
- C. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.
- D. Create an IAM policy than allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

ANSWER: D

Explanation:

Create an IAM policy than allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account. is correct because the project account is independent and the requirement must be enforced inside that account. An identity-based IAM policy attached to the IAM roles or groups used by developers can control the EC2 RunInstances action and constrain it with policy condition keys. For EC2, policies can use

condition keys such as `ec2:InstanceType` to permit only a specific instance type, such as `t3.small`. AWS global condition keys such as `aws:RequestedRegion` can restrict where the request is made, allowing launches only in `us-east-2`. In practice, the policy should allow `RunInstances` only when both conditions match, and organizations often add an explicit deny guardrail for noncompliant instance types or Regions to make the restriction harder to bypass through other permissions. This approach directly satisfies the requirement without needing AWS Organizations and applies consistently to the developers' normal access paths. See the Amazon EC2 service authorization reference for EC2 condition keys at [Amazon EC2 actions, resources, and condition keys](#) and the IAM documentation for [aws:RequestedRegion](#).

QUESTION NO: 22

A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with the number of files rapidly declining after business hours.

What is the MOST cost-effective migration recommendation?

- A.** Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in an Amazon S3 bucket.
- B.** Create a queue using Amazon MQ. Configure the existing web server to publish to the new queue. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the files. Store the processed files in Amazon EFS. Shut down the EC2 instance after the task is complete.
- C.** Create a queue using Amazon MQ. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in Amazon EFS.
- D.** Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the files. Scale the EC2 instances based on the SQS queue length. Store the processed files in an Amazon S3 bucket.

ANSWER: D

Explanation:

Using Amazon SQS with Amazon EC2 instances in an EC2 Auto Scaling group and storing processed files in Amazon S3 is the most cost-effective fit for this workload. The queue decouples file ingestion from processing, so the web tier can accept uploads quickly while processing capacity scales independently as backlog changes. Because each media file can take up to 1 hour to process, EC2 worker instances are appropriate for the long-running processing jobs, and Auto Scaling can add or remove workers based on the SQS queue depth or backlog-per-instance metrics. This matches the described demand pattern: more instances can run during business hours when the queue grows, and capacity can scale down after hours as the backlog declines, reducing compute cost.

Amazon S3 is also a cost-effective and durable destination for processed media files, especially when replacing NAS-style storage for object-based file outputs. AWS documents this queue-based scaling pattern for EC2 Auto Scaling with SQS, including scaling based on queue backlog, and S3 is designed for highly durable object storage. References: [Scaling based on Amazon SQS](#) and [Amazon S3 User Guide](#).

QUESTION NO: 23

A company uses AWS Organizations to manage more than 1,000 AWS accounts. The company has created a new developer organization. There are 540 developer member accounts that must be moved to the new developer organization. All accounts are set up with all the required information so that each account can be operated as a standalone account.

Which combination of steps should a solutions architect take to move all of the developer accounts to the new developer organization? (Select THREE.)

- A. Call the MoveAccount operation in the Organizations API from the old organization ' s management account to migrate the developer accounts to the new developer organization.
- B. From the management account, remove each developer account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.
- C. From each developer account, remove the account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.
- D. Sign in to the new developer organization ' s management account and create a placeholder member account that acts as a target for the developer account migration.
- E. Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization ' s management account to send invitations to the developer accounts.
- F. Have each developer sign in to their account and confirm to join the new developer organization.

ANSWER: B E F

Explanation:

The correct sequence is to first detach the developer accounts from the existing organization, then invite them into the new developer organization, and finally have each account accept the invitation. From the management account, remove each developer account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API is correct because this API operation is performed by the current organization's management account and removes a member account so that it becomes a standalone AWS account. This is why the question states that all required standalone account information is already configured. Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization's management account to send invitations to the developer accounts is also correct because the destination organization's management account must initiate the invitation process. Have each developer sign in to their account and confirm to join the new developer organization is correct because an invited standalone account must accept the organization handshake before it becomes a member of the new organization. AWS documents these behaviors in the Organizations API references for [RemoveAccountFromOrganization](#) and [InviteAccountToOrganization](#).

QUESTION NO: 24

A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts.

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

- A. In the centralized account, create an IAM role that has the Lambda service as a trusted entity. Add an inline policy to assume the roles of the other AWS accounts.
- B. In the other AWS accounts, create an IAM role that has minimal permissions. Add the centralized account's Lambda IAM role as a trusted entity.
- C. In the centralized account, create an IAM role that has roles of the other accounts as trusted entities. Provide minimal permissions.
- D. In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized account. Add the Lambda service as a trusted entity.
- E. In the other AWS accounts, create an IAM role that has minimal permissions. Add the Lambda service as a trusted entity.

ANSWER: A B

Explanation:

The correct combination is to configure the Lambda execution role in the centralized account so that AWS Lambda can assume it, and then allow that role to assume narrowly scoped roles in the other AWS accounts. "In the centralized account,

create an IAM role that has the Lambda service as a trusted entity. Add an inline policy to assume the roles of the other AWS accounts.” is correct because Lambda needs an execution role whose trust policy allows the Lambda service principal to use the role, and whose permissions policy allows only the required sts:AssumeRole actions against specific target role ARNs.

“In the other AWS accounts, create an IAM role that has minimal permissions. Add the centralized account's Lambda IAM role as a trusted entity.” is also correct because each target account should expose a role with only the permissions needed for governance and remediation in that account. Its trust policy should name the centralized Lambda execution role as the trusted principal. This is the standard cross-account access pattern and supports least privilege by separating who can assume the role from what the role can do after assumption. See AWS guidance on [cross-account IAM roles](#) and [Lambda execution roles](#).

QUESTION NO: 25

A company is creating a centralized logging service running on Amazon EC2 that will receive and analyze logs from hundreds of AWS accounts. AWS PrivateLink is being used to provide connectivity between the client services and the logging service.

In each AWS account with a client, an interface endpoint has been created for the logging service and is available. The logging service running on EC2 instances with a Network Load Balancer (NLB) are deployed in different subnets. The clients are unable to submit logs using the VPC endpoint.

Which combination of steps should a solutions architect take to resolve this issue? (Select TWO.)

- A. Check that the NACL is attached to the logging service subnet to allow communications to and from the NLB subnets. Check that the NACL is attached to the NLB subnet to allow communications to and from the logging service subnets running on EC2 instances.
- B. Check that the NACL is attached to the logging service subnets to allow communications to and from the interface endpoint subnets. Check that the NACL is attached to the interface endpoint subnet to allow communications to and from the logging service subnets running on EC2 instances.
- C. Check the security group for the logging service running on the EC2 instances to ensure it allows Ingress from the NLB subnets.
- D. Check the security group for the logging service running on EC2 instances to ensure it allows ingress from the clients.
- E. Check the security group for the NLB to ensure it allows ingress from the interface endpoint subnets.

ANSWER: A C

Explanation:

The correct remediation is to validate the network controls between the Network Load Balancer subnets and the EC2 subnets that host the logging service, and to ensure the EC2 security group permits traffic from the Network Load Balancer subnets. With AWS PrivateLink endpoint services, consumers connect to an interface endpoint in their own VPC, and that traffic is delivered to the provider service through a Network Load Balancer. For endpoint service traffic, the backend targets do not normally see the original client IP as the source; the traffic reaching the targets is associated with the Network Load Balancer nodes. Therefore, the EC2 instances that run the logging service must allow inbound traffic from the Network Load Balancer subnets or node addresses. Because network ACLs are stateless and apply at the subnet boundary, the ACLs associated with both the Network Load Balancer subnets and the EC2 logging service subnets must allow the required inbound and outbound flows, including ephemeral response traffic. AWS documents this PrivateLink endpoint service architecture and the use of Network Load Balancers for endpoint services in [AWS PrivateLink endpoint service documentation](#), and describes Network Load Balancer target connectivity behavior in the [Network Load Balancer target group documentation](#).

QUESTION NO: 26

A solutions architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from

upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.

Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

- A.** Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
- B.** Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
- C.** Modify the application to use Amazon DynamoDB instead of Amazon RDS. Configure Auto Scaling for the DynamoDB table. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
- D.** Update the application to use a Redis task queue instead of the in-memory queue. Build a Docker container image for the application. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

ANSWER: B

Explanation:

The approach that creates an API Gateway REST API with a service proxy to place requests on an Amazon SQS queue, then processes the queue with smaller EC2 instances in an Auto Scaling group, is the best fit. API Gateway can quickly accept the unchanged HTTP requests and persist each work item to SQS, allowing clients to receive a timely success response without waiting for the 90-second backend processing step. Amazon SQS provides durable, highly available message storage, which removes the data-loss risk of an in-memory queue on a single instance and cleanly decouples request ingestion from processing. The processing tier can then scale horizontally because the application already supports parallel work. Scaling the Auto Scaling group based on the approximate number of messages in the queue directly matches capacity to backlog size, reducing processing latency during bursts while allowing the fleet to shrink during quiet periods to control cost. This is a standard AWS pattern for asynchronous, resilient workload processing. See the Amazon SQS documentation on reliable, scalable queues at [Amazon SQS Developer Guide](#) and the guidance for scaling EC2 Auto Scaling groups using SQS metrics at [Scaling based on Amazon SQS](#).

QUESTION NO: 27

A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet. The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 TB of data from an S3 bucket each day. The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations. Which solution will meet these requirements?

- A.** Replace the NAT gateways with NAT instances. In the VPC route table, create a route from the private subnets to the NAT instances.
- B.** Move the EC2 instances to the public subnets. Remove the NAT gateways.

- C. Set up an S3 gateway VPC endpoint in the VPAttach an endpoint policy to the endpoint to allow the required actions on the S3 bucket.
- D. Attach an Amazon Elastic File System (Amazon EFS) volume to the EC2 instances. Host the images on the EFS volume.

ANSWER: C

Explanation:

Set up an S3 gateway VPC endpoint in the VPAttach an endpoint policy to the endpoint to allow the required actions on the S3 bucket is correct because it keeps the EC2 instances in private subnets while allowing private, highly available access to Amazon S3 without sending S3 traffic through NAT gateways. Gateway endpoints for Amazon S3 are horizontally scaled, redundant, and do not require ongoing instance management, so this approach preserves the operational simplicity of the current architecture. It also reduces cost because traffic from the private subnets to S3 can use the gateway endpoint route instead of incurring NAT gateway data processing charges for approximately 1 TB of daily S3 retrievals. Security is maintained or improved because access to S3 can be constrained with endpoint policies, bucket policies, and IAM permissions, and the instances do not need public IP addresses or direct internet exposure for S3 access. AWS specifically recommends gateway endpoints as a way to access S3 from a VPC without an internet gateway or NAT device. See the AWS documentation for [Gateway endpoints for Amazon S3](#) and [controlling access to VPC endpoints using endpoint policies](#).

QUESTION NO: 28

A solutions architect has launched multiple Amazon EC2 instances in a placement group within a single Availability Zone. Because of additional load on the system, the solutions architect attempts to add new instances to the placement group. However, the solutions architect receives an insufficient capacity error.

What should the solutions architect do to troubleshoot this issue?

- A. Use a spread placement group. Set a minimum of eight instances for each Availability Zone.
- B. Stop and start all the instances in the placement group. Try the launch again.
- C. Create a new placement group. Merge the new placement group with the original placement group.
- D. Launch the additional instances as Dedicated Hosts in the placement groups.

ANSWER: B

Explanation:

Stop and start all the instances in the placement group. Try the launch again. is the correct action because this scenario matches AWS guidance for capacity errors in an existing cluster placement group. Cluster placement groups place instances physically close together inside a single Availability Zone to provide low-latency, high-throughput networking. Because that placement depends on available contiguous capacity, adding more instances later can fail if AWS cannot fit the requested capacity near the already-running instances. AWS recommends launching all needed instances into a cluster placement group at the same time when possible. If an insufficient capacity error occurs after instances are already running, stopping and starting all instances in the placement group gives AWS an opportunity to place the entire set of instances together on hardware that has enough capacity for the group. After the restart, the architect can retry launching the additional instances. This approach directly addresses the placement and capacity constraint while preserving the intent of using the placement group for tightly coupled workloads. See the AWS documentation for placement group behavior and limitations in [Amazon EC2 placement groups](#) and general instance stop/start behavior in [Stop and start Amazon EC2 instances](#).

QUESTION NO: 29

A company wants to use AWS to create a business continuity solution in case the company's main on-premises application fails. The application runs on physical servers that also run other applications. The on-premises application that the company is planning to migrate uses a MySQL database as a data store. All the company's on-premises applications use operating systems that are compatible with Amazon EC2. Which solution will achieve the company's goal with the LEAST operational overhead?

A. Install the AWS Replication Agent on the source servers, including the MySQL servers. Set up replication for all servers. Launch test instances for regular drills. Cut over to the test instances to fail over the workload in the case of a failure event.

B. Install the AWS Replication Agent on the source servers, including the MySQL servers. Initialize AWS Elastic Disaster Recovery in the target AWS Region. Define the launch settings. Frequently perform failover and fallback from the most recent point in time.

C. Create AWS Database Migration Service (AWS DMS) replication servers and a target Amazon Aurora MySQL DB cluster to host the database. Create a DMS replication task to copy the existing data to the target DB cluster. Create a local AWS Schema Conversion Tool (AWS SCT) change data capture (CDC) task to keep the data synchronized. Install the rest of the software on EC2 instances by starting with a compatible base AMI.

D. Deploy an AWS Storage Gateway Volume Gateway on premises. Mount volumes on all on- premises servers. Install the application and the MySQL database on the new volumes. Take regular snapshots. Install all the software on EC2 Instances by starting with a compatible base AMI. Launch a Volume Gateway on an EC2 instance. Restore the volumes from the latest snapshot. Mount the new volumes on the EC2 instances in the case of a failure event.

ANSWER: B

Explanation:

Install the AWS Replication Agent on the source servers, including the MySQL servers. Initialize AWS Elastic Disaster Recovery in the target AWS Region. Define the launch settings. Frequently perform failover and fallback from the most recent point in time is correct because AWS Elastic Disaster Recovery is the managed AWS service designed specifically for low-overhead business continuity and disaster recovery of physical, virtual, and cloud-based servers into AWS. After the AWS Replication Agent is installed, the service continuously replicates block-level data from the source servers to a low-cost staging area in the selected AWS Region. Recovery instances can then be launched on Amazon EC2 using predefined launch settings, which reduces the amount of manual rebuild, reconfiguration, and database migration work required during an outage.

This approach fits the scenario because the source systems use operating systems compatible with Amazon EC2 and include a MySQL server that can be replicated at the server level along with the application environment. AWS Elastic Disaster Recovery supports non-disruptive drills, point-in-time recovery, failover, and fallback workflows, making it well suited for maintaining a ready recovery environment with minimal ongoing operational effort. See [AWS Elastic Disaster Recovery documentation](#) and [AWS Disaster Recovery on AWS](#).

QUESTION NO: 30

A company operates a proxy server on a fleet of Amazon EC2 instances. Partners in different countries use the proxy server to test the company's functionality. The EC2 instances are running in a VPC. and the instances have access to the internet.

The company's security policy requires that partners can access resources only from domains that the company owns.

Which solution will meet these requirements?

A. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. Configure a rule that has a low numeric value that allows requests for domains in the allowed list. Associate the rule group with the VPC.

B. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a Route 53 outbound endpoint. Associate the outbound endpoint with the VPC. Associate the domain list with the outbound endpoint.

C. Create an Amazon Route 53 traffic flow policy to match the allowed domains. Configure the traffic flow policy to forward requests that match to the Route 53 Resolver. Associate the traffic flow policy with the VPC.

D. Create an Amazon Route 53 outbound endpoint. Associate the outbound endpoint with the VPC. Configure a Route 53 traffic flow policy to forward requests for allowed domains to the outbound endpoint. Associate the traffic flow policy with the VPC.

ANSWER: A

Explanation:

Creating an Amazon Route 53 Resolver DNS Firewall domain list that contains the company-owned allowed domains is the correct solution because DNS Firewall is designed to control outbound DNS resolution for resources in a VPC. By associating a DNS Firewall rule group with the VPC, DNS queries from the EC2-based proxy instances can be evaluated before the instances connect to external resources. A lower numeric priority rule that allows queries for the approved domain list is evaluated first, and a higher numeric priority rule that blocks all other DNS requests enforces a default-deny posture. This matches the requirement that partners using the proxy can reach only resources in domains owned by the company.

Route 53 Resolver DNS Firewall supports rule actions such as ALLOW and BLOCK and evaluates rules in priority order, where lower values are processed before higher values. Associating the rule group with the VPC applies the filtering behavior to DNS queries that originate from that VPC. For more details, see [Amazon Route 53 Resolver DNS Firewall](#) and [DNS Firewall rule groups and rules](#).

QUESTION NO: 31

A company has a transit gateway that connects multiple VPCs in the same AWS Region. The company needs a centralized way to inspect network traffic and allow internet access for the workload VPCs.

Which solution meets these requirements?

- A. Create a Gateway Load Balancer (GWLB), GWLB endpoints, and a network virtual appliance in an existing workload VPC
- B. Create a Gateway Load Balancer (GWLB) in an existing workload VPC
- C. Create an inspection VPC and an internet access VPC
- D. Create an inspection VPC that contains a Gateway Load Balancer (GWLB), GWLB endpoints, and a network virtual appliance. Update the route tables in all workload VPCs to send traffic to the transit gateway. Configure the transit gateway route tables to forward traffic to the GWLB endpoints. Enable appliance mode on the transit gateway.
- E. Create a centralized inspection/egress VPC with a Gateway Load Balancer, GWLB endpoints, network virtual appliances, NAT gateways, and an internet gateway. Route workload VPC internet-bound traffic to the transit gateway, route the transit gateway to the inspection VPC attachment, route inspection VPC traffic through the GWLB endpoints, and enable appliance mode.

ANSWER: E

Explanation:

Create a centralized inspection/egress VPC with a Gateway Load Balancer, GWLB endpoints, network virtual appliances, NAT gateways, and an internet gateway is correct because it implements the standard AWS pattern for centralized traffic inspection and shared internet egress. Workload VPCs send internet-bound traffic to the transit gateway, and the transit gateway routes that traffic to the inspection VPC attachment. Inside the inspection VPC, subnet route tables steer traffic through Gateway Load Balancer endpoints so that the security appliances can inspect traffic transparently and scale horizontally behind the Gateway Load Balancer. After inspection, outbound traffic can use NAT gateways and the internet gateway to reach the internet. Enabling appliance mode on the transit gateway VPC attachment is important for stateful inspection because it helps keep forward and return flows symmetrically pinned to the same Availability Zone and appliance path. This design aligns with AWS guidance for deploying centralized inspection appliances with Transit Gateway and Gateway Load Balancer. See [AWS Transit Gateway appliance scenarios](#) and [Gateway Load Balancer documentation](#).

QUESTION NO: 32

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 Elastic Kubernetes Service (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 Database Migration Service (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

Explanation:

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. is the correct solution because AWS Migration Evaluator is designed specifically to create a data-driven migration business case and total cost of ownership estimate before moving workloads to AWS. The collector gathers on-premises inventory and utilization data, which is important for sizing AWS resources realistically instead of relying only on static server specifications. After the data is imported, a scenario can be configured to model the target AWS environment and produce cost projections for the migrated workload. For this case, the assessment can use the current Kubernetes infrastructure and database footprint to help estimate the AWS run-rate for an Amazon EKS deployment with managed node groups and an Amazon RDS for PostgreSQL database. The Quick Insights report is intended to provide a summarized business-case view, including estimated AWS costs and savings opportunities, which matches the requirement to obtain TCO information before migration. AWS describes Migration Evaluator as a service for building a directional business case for AWS migration, and its reports are commonly used during early planning and assessment phases. See [AWS Migration Evaluator](#) and [What is Migration Evaluator?](#).

QUESTION NO: 33

A company wants to migrate to AWS. The company wants to use a multi-account structure with centrally managed access to all accounts and applications. The company also wants to keep the traffic on a private network. Multi-factor authentication (MFA) is required at login, and specific roles are assigned to user groups.

The company must create separate accounts for development, staging, production, and shared network. The production account and the shared network account must have connectivity to all accounts. The development account and the staging account must have access only to each other.

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

A. Deploy a landing zone environment by using AWS Control Tower. Enroll accounts and invite existing accounts into the resulting organization in AWS Organizations.

B. Enable AWS Security Hub in all accounts to manage cross-account access. Collect findings through AWS CloudTrail to force MFA login.

C. Create transit gateways and transit gateway VPC attachments in each account. Configure appropriate route tables.

D. Set up and enable AWS IAM Identity Center (AWS Single Sign-On). Create appropriate permission sets with required MFA for existing accounts.

E. Enable AWS Control Tower in all Recounts to manage routing between accounts. Collect findings through AWS CloudTrail to force MFA login.

F. Create IAM users and groups. Configure MFA for all users. Set up Amazon Cognito user pools and identity pools to manage access to accounts and between accounts.

ANSWER: A C D

Explanation:

Deploy a landing zone environment by using AWS Control Tower. Enroll accounts and invite existing accounts into the resulting organization in AWS Organizations is correct because AWS Control Tower provides the recommended managed landing zone for a governed multi-account AWS environment, including account enrollment, organizational units, and

centralized guardrails. Set up and enable AWS IAM Identity Center (AWS Single Sign-On). Create appropriate permission sets with required MFA for existing accounts is correct because IAM Identity Center is designed for centralized workforce access across multiple AWS accounts, mapping users and groups to permission sets while enforcing MFA and role-based access. Create transit gateways and transit gateway VPC attachments in each account. Configure appropriate route tables is correct because AWS Transit Gateway provides private, scalable VPC-to-VPC connectivity across accounts. Separate transit gateway route tables and associations/propagations can implement the required connectivity pattern, where production and shared network can reach all accounts while development and staging can be limited to each other. See [AWS Control Tower documentation](#) and [AWS Transit Gateway documentation](#).

QUESTION NO: 34

A company has applications in an AWS account that is named Source. The account is in an organization in AWS Organizations. One of the applications uses AWS Lambda functions and store's inventory data in an Amazon Aurora database. The application deploys the Lambda functions by using a deployment package. The company has configured automated backups for Aurora.

The company wants to migrate the Lambda functions and the Aurora database to a new AWS account that is named Target. The application processes critical data, so the company must minimize downtime.

Which solution will meet these requirements?

- A.** Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the automated Aurora DB cluster snapshot with the Target account.
- B.** Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the Aurora DB cluster with the Target account by using AWS Resource Access Manager (AWS RAM). Grant the Target account permission to clone the Aurora DB cluster.
- C.** Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions and the Aurora DB cluster with the Target account. Grant the Target account permission to clone the Aurora DB cluster.
- D.** Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions with the Target account. Share the automated Aurora DB cluster snapshot with the Target account.

ANSWER: B

Explanation:

Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the Aurora DB cluster with the Target account by using AWS Resource Access Manager (AWS RAM). Grant the Target account permission to clone the Aurora DB cluster. is correct because it combines a practical Lambda migration method with the lowest-downtime Aurora migration mechanism available in this scenario. For Lambda functions that use a deployment package, the function code can be retrieved from the source function and used to create equivalent functions in the new account, while the surrounding configuration, IAM roles, triggers, and environment settings can be recreated as part of the migration. For the database, Aurora fast cloning uses a copy-on-write approach, so creating the target cluster clone does not require a full snapshot copy and restore cycle. AWS supports sharing Aurora DB clusters with another AWS account through AWS Resource Access Manager so that the target account can create a clone of the shared cluster. This approach is especially suitable for critical applications because the database copy can be created quickly and the required downtime can be limited mainly to the final application cutover. See the AWS documentation for [Aurora cloning](#) and [Lambda deployment packages](#).

QUESTION NO: 35

A company's security compliance requirements state that all Amazon EC2 images must be scanned for vulnerabilities and must pass a CVE assessment. A solutions architect is developing a mechanism to create security-approved AMIs that can be used by developers. Any new AMIs should go through an automated assessment process and be marked as approved before developers can use them. The approved images must be scanned every 30 days to ensure compliance.

Which combination of steps should the solutions architect take to meet these requirements while following best practices? (Select TWO)

- A. Use the AWS Systems Manager EC2 agent to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned
- B. Use AWS Lambda to write automatic approval rules Store the approved AMI list in AWS Systems Manager Parameter Store Use Amazon EventBridge to trigger an AWS Systems Manager Automation document on all EC2 instances every 30 days.
- C. Use Amazon Inspector to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned
- D. Use AWS Lambda to write automatic approval rules Store the approved AMI list in AWS Systems Manager Parameter Store Use a managed AWS Config rule for continuous scanning on all EC2 instances, and use AWS Systems Manager Automation documents for remediation
- E. Use AWS CloudTrail to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned

ANSWER: B C

Explanation:

Use Amazon Inspector to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned is correct because Amazon Inspector is the AWS-native vulnerability management service for Amazon EC2 workloads. It evaluates instances for software vulnerabilities and produces findings that include CVE information, which fits the requirement to assess AMIs for vulnerabilities by launching test instances from those AMIs and scanning them. Use AWS Lambda to write automatic approval rules Store the approved AMI list in AWS Systems Manager Parameter Store Use Amazon EventBridge to trigger an AWS Systems Manager Automation document on all EC2 instances every 30 days. is also correct because it provides an automated approval and recurring validation workflow. Lambda can evaluate Inspector findings and update an approved AMI list in Systems Manager Parameter Store, giving developers a controlled source of approved image IDs. EventBridge can schedule the recurring 30-day process, and Systems Manager Automation can orchestrate repeatable operational tasks such as launching or targeting instances for validation and enforcing the scan workflow. Together, these services implement automated assessment, approval tracking, and periodic compliance revalidation using AWS managed capabilities. See [Amazon Inspector EC2 instance scanning](#) and [AWS Systems Manager Automation](#).

QUESTION NO: 36

A company that is developing a mobile game is making game assets available in two AWS Regions. Game assets are served from a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in each Region. The company requires game assets to be (etched from the closest Region. If game assets become unavailable in the closest Region, they should be fetched from the other Region.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution. Create an origin group with one origin for each ALB. Set one of the origins as primary.
- B. Create an Amazon Route 53 health check for each ALB. Create a Route 53 failover routing record pointing to the two ALBs. Set the Evaluate Target Health value to Yes.
- C. Create two Amazon CloudFront distributions, each with one ALB as the origin. Create an Amazon Route 53 failover routing record pointing to the two CloudFront distributions. Set the Evaluate Target Health value to Yes.
- D. Create an Amazon Route 53 health check for each ALB. Create a Route 53 latency alias record pointing to the two ALBs. Set the Evaluate Target Health value to Yes.

ANSWER: D

Explanation:

Create an Amazon Route 53 health check for each ALB. Create a Route 53 latency alias record pointing to the two ALBs. Set the Evaluate Target Health value to Yes is correct because Route 53 latency-based routing is designed for workloads deployed in multiple AWS Regions where users should be sent to the Region that provides the lowest network latency. By

creating latency alias records for the regional Application Load Balancers, Route 53 can choose the best regional endpoint for each user request based on measured latency. Enabling health evaluation ensures Route 53 considers endpoint availability when answering DNS queries, so if the normally preferred regional ALB or its targets are unhealthy, Route 53 can stop returning that endpoint and return the healthy regional endpoint instead. This matches the requirement to fetch assets from the closest Region during normal operation and automatically use the other Region if the closest Region becomes unavailable. AWS documents latency routing as the appropriate policy when resources are in multiple AWS Regions and traffic should be routed to the Region with the best latency. Alias records to Elastic Load Balancing endpoints are also a standard Route 53 pattern for highly available DNS routing. References: [AWS Route 53 latency-based routing](#) and [AWS Route 53 alias record values](#).

QUESTION NO: 37

A company is running an application in the AWS Cloud. The company has several third-party services that integrate with the application through a RESTful API. The API is a serverless implementation with an Amazon API Gateway regional API endpoint that integrates with several different AWS Lambda functions.

The application's data is nonrelational and is stored in an Amazon DynamoDB table. The application and the API are running in the eu-west-1 Region. The company needs the API to also be available in the us-east-1 Region. All data must be available in both Regions. A solutions architect already has deployed all the Lambda functions in us-east-1

Which additional steps should the solutions architect take to meet these requirements? (Select TWO.)

- A. Deploy a second API Gateway regional API endpoint in us-east-1. Create Lambda integration with the functions in us-east-1.
- B. Enable DynamoDB Streams on the table in eu-west-1. Replicate all changes to a DynamoDB table in us-east-1
- C. Modify the DynamoDB table to be a global table in eu-west-1 and in us-east-1.
- D. Change the API Gateway API endpoint in eu-west-1 to an edge-optimized endpoint. Create Lambda integration with the functions in both Regions.
- E. Create a DynamoDB read replica in us-east-1.

ANSWER: A C

Explanation:

Deploy a second API Gateway regional API endpoint in us-east-1. Create Lambda integration with the functions in us-east-1. is correct because a Regional API Gateway endpoint is deployed in a specific AWS Region and is intended to serve clients from that Region while integrating with regional backend resources such as Lambda functions. Since the Lambda functions already exist in us-east-1, creating a separate Regional API endpoint there and integrating it with those functions makes the API available from both eu-west-1 and us-east-1 using a regional, serverless architecture. Modify the DynamoDB table to be a global table in eu-west-1 and in us-east-1. is also correct because DynamoDB global tables provide fully managed, multi-Region, active-active replication. This ensures that the application's nonrelational data is available in both Regions and that writes in either Region can be replicated automatically to the other Region. Together, these steps provide regional API availability and multi-Region data access without requiring custom replication logic. For more information, see the AWS documentation for [API Gateway endpoint types](#) and [DynamoDB global tables](#).

QUESTION NO: 38

A company is currently in the design phase of an application that will need an RPO of less than 5 minutes and an RTO of less than 10 minutes. The solutions architecture team is forecasting that the database will store approximately 10 TB of data. As part of the design, they are looking for a database solution that will provide the company with the ability to fail over to a secondary Region. Which solution will meet these business requirements at the LOWEST cost?

- A. Deploy an Amazon Aurora DB cluster and take snapshots of the cluster every 5 minutes. Once a snapshot is complete, copy the snapshot to a secondary Region to serve as a backup in the event of a failure.
- B. Deploy an Amazon RDS instance with a cross-Region read replica in a secondary Region. In the event of a failure, promote the read replica to become the primary.

C. Deploy an Amazon Aurora DB cluster in the primary Region and another in a secondary Region. Use AWS DMS to keep the secondary Region in sync.

D. Deploy an Amazon RDS instance with a read replica in the same Region. In the event of a failure, promote the read replica to become the primary.

ANSWER: B

Explanation:

Deploy an Amazon RDS instance with a cross-Region read replica in a secondary Region. In the event of a failure, promote the read replica to become the primary is the correct choice because Amazon RDS cross-Region read replicas are designed for disaster recovery scenarios where a standby database must be available in another AWS Region. Replication is asynchronous, and with the database instance class, storage, networking, and workload sized correctly, replica lag can typically be kept low enough to satisfy an RPO of less than 5 minutes. If the primary Region becomes unavailable, the read replica can be promoted to a standalone DB instance, allowing the application to resume writes in the secondary Region within the required recovery window when paired with appropriate application failover or DNS cutover procedures. This approach is also cost-effective because it uses a single primary database and one replica rather than maintaining a more complex active-active or separately synchronized database architecture. AWS documents cross-Region read replicas as a way to improve disaster recovery capabilities, and also documents promotion of a read replica to a standalone DB instance. References: [Amazon RDS cross-Region read replicas](#) and [Promoting a read replica to be a standalone DB instance](#).

QUESTION NO: 39

An enterprise company is building an infrastructure services platform for its users. The company has the following requirements:

Provide least privilege access to users when launching AWS infrastructure so users cannot provision unapproved services.

Use a central account to manage the creation of infrastructure services.

Provide the ability to distribute infrastructure services to multiple accounts in AWS Organizations.

Provide the ability to enforce tags on any infrastructure that is started by users.

Which combination of actions using AWS services will meet these requirements? (Choose three.)

A. Develop infrastructure services using AWS Cloud Formation templates. Add the templates to a central Amazon S3 bucket and add the IAM roles or users that require access to the S3 bucket policy.

B. Develop infrastructure services using AWS Cloud Formation templates. Upload each template as an AWS Service Catalog product to portfolios created in a central AWS account. Share these portfolios with the Organizations structure created for the company.

C. Allow user IAM roles to have `AWSCloudFormationFullAccess` and `AmazonS3ReadOnlyAccess` permissions. Add an Organizations SCP at the AWS account root user level to deny all services except AWS CloudFormation and Amazon S3.

D. Allow user IAM roles to have `ServiceCatalogEndUserAccess` permissions only. Use an automation script to import the central portfolios to local AWS accounts, copy the `TagOption` assign users access and apply launch constraints.

E. Use the AWS Service Catalog `TagOption` Library to maintain a list of tags required by the company. Apply the `TagOption` to AWS Service Catalog products or portfolios.

F. Use the AWS CloudFormation `Resource Tags` property to enforce the application of tags to any CloudFormation templates that will be created for users.

ANSWER: B D E

Explanation:

AWS Service Catalog is the right fit because it lets a central platform team define approved infrastructure as products, package those products from AWS CloudFormation templates, and publish them through portfolios. "Develop infrastructure

services using AWS Cloud Formation templates. Upload each template as an AWS Service Catalog product to portfolios created in a central AWS account. Share these portfolios with the Organizations structure created for the company.” satisfies the central management and multi-account distribution requirements by using portfolio sharing with AWS Organizations. “Allow user IAM roles to have ServiceCatalogEndUserAccess permissions only. Use an automation script to import the central portfolios to local AWS accounts, copy the TagOption assign users access and apply launch constraints.” supports least privilege because end users interact with approved Service Catalog products rather than receiving broad permissions to create arbitrary AWS resources. Launch constraints allow provisioning to occur through controlled IAM roles. “Use the AWS Service Catalog TagOption Library to maintain a list of tags required by the company. Apply the TagOption to AWS Service Catalog products or portfolios.” meets the tagging requirement because TagOptions define allowed tag keys and values that can be associated with products and portfolios. See the AWS documentation for [AWS Service Catalog](#) and [TagOptions](#).

QUESTION NO: 40

A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalogue page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.

Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times.

Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Select TWO.)

- A.** Configure read replicas for Amazon RDS MySQL and use the single reader endpoint in the web application to reduce the load on the backend database tier.
- B.** Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionality. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
- C.** Configure the target group health check to use a TCP check of the Amazon EC2 web server and the Amazon Route 53 health check against the product page to evaluate full application functionality. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
- D.** Configure an Amazon CloudWatch alarm for Amazon RDS with an action to recover a high-load, impaired RDS instance in the database tier.
- E.** Configure an Amazon ElastiCache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

ANSWER: B E

Explanation:

Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionality. Configure Amazon CloudWatch alarms to notify administrators when the site fails. is correct because the ALB target group health check should validate whether each web server can receive and respond to traffic, not whether every downstream dependency is fast and healthy. Using a lightweight local page prevents Auto Scaling from unnecessarily replacing healthy EC2 instances when the database is slow, while a Route 53 health check against the product page provides an external, end-to-end signal for the full application path and can publish metrics for CloudWatch alarms. See [Amazon Route 53 health check types](#).

Configure an Amazon ElastiCache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

is also correct because a product catalog is commonly read-heavy and well suited to caching. ElastiCache can absorb repeated reads, reduce query volume against RDS MySQL, lower database latency, and provide a scalable layer for future growth. AWS documents this pattern as a way to improve application performance and reduce pressure on backend databases: [ElastiCache caching strategies](#).

QUESTION NO: 41

A software development company has multiple engineers who are working remotely. The company is running Active Directory Domain Services (AD DS) on an Amazon EC2 instance. The company's security policy states that all internal, nonpublic services that are deployed in a VPC must be accessible through a VPN. Multi-factor authentication (MFA) must be used for access to a VPN.

What should a solution architect do to meet these requirements?

- A.** Create an AWS Site-to-Site VPN connection. Configure integration between a VPN and AD DS. Use an Amazon Workspaces client with MFA support enabled to establish a VPN connection.
- B.** Create an AWS Client VPN endpoint. Create an AD Connector directory for integration with AD DS. Enable MFA for AD Connector. Use AWS Client VPN to establish a VPN connection.
- C.** Create multiple AWS Site-to-Site VPN connections by using AWS VPN CloudHub. Configure integration between AWS VPN CloudHub and AD DS. Use AWS Cop4ot to establish a VPN connection.
- D.** Create an Amazon WorkLink endpoint. Configure integration between Amazon WorkLink and AD DS. Enable MFA in Amazon WorkLink. Use AWS Client VPN to establish a VPN connection.

ANSWER: B

Explanation:

Create an AWS Client VPN endpoint. Create an AD Connector directory for integration with AD DS. Enable MFA for AD Connector. Use AWS Client VPN to establish a VPN connection is correct because AWS Client VPN is designed to provide secure remote-user VPN access into AWS VPCs. It lets individual remote engineers connect to private, nonpublic resources in a VPC over a managed VPN endpoint, which matches the requirement for remote access to internal services. AWS Client VPN supports authentication through AWS Directory Service, including AD Connector, which can connect to an existing self-managed Microsoft Active Directory environment such as AD DS running on Amazon EC2. By creating an AD Connector directory and associating it with the Client VPN endpoint, users can authenticate with existing domain credentials. MFA can be enabled for AD Connector through RADIUS-based MFA, satisfying the security policy that VPN access must require multi-factor authentication. This approach avoids building and operating custom VPN infrastructure while using AWS-managed Client VPN for scalable remote access. Relevant AWS documentation includes [What is AWS Client VPN?](#) and [Enable multi-factor authentication for AD Connector](#).

QUESTION NO: 42

A company is migrating its infrastructure to the AWS Cloud. The company must comply with a variety of regulatory standards for different projects. The company needs a multi-account environment.

A solutions architect needs to prepare the baseline infrastructure. The solution must provide a consistent baseline of management and security, but it must allow flexibility for different compliance requirements within various AWS accounts. The solution also needs to integrate with the existing on-premises Active Directory Federation Services (AD FS) server.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A.** Create an organization in AWS Organizations. Create a single SCP for least privilege access across all accounts. Create a single OU for all accounts. Configure an IAM identity provider for federation with the on-premises AD FS server. Configure a central logging account with a defined process for log generating services to send log events to the central account. Enable AWS Config in the central account with conformance packs for all accounts.
- B.** Create an organization in AWS Organizations. Enable AWS Control Tower on the organization. Review included controls (guardrails) for SCPs. Check AWSConfig for areas that require additions. Add OUS as necessary. Connect AWS IAM Identity Center (AWS Single Sign-On) to the on-premises AD FS server.
- C.** Create an organization in AWS Organizations. Create SCPs for least privilege access. Create an OU structure, and use it to group AWS accounts. Connect AWS IAM Identity Center (AWS Single Sign-On) to the on-premises AD FS server. Configure a central logging account with a defined process for log generating services to send log events to the central account. Enable AWS Config in the central account with aggregators and conformance packs.

D. Create an organization in AWS Organizations. Enable AWS Control Tower on the organization. Review included controls (guardrails) for SCPs. Check AWSConfig for areas that require additions. Configure an IAM identity provider for federation with the on-premises AD FS server.

ANSWER: B

Explanation:

The solution that creates an organization in AWS Organizations, enables AWS Control Tower, reviews included controls, adds organizational units as necessary, and connects AWS IAM Identity Center to the on-premises AD FS server is correct because it uses AWS-managed capabilities to establish and govern a multi-account landing zone with minimal custom build effort. AWS Control Tower automates the creation of a secure baseline environment, including a management account structure, core accounts, preventive and detective controls, centralized logging, and governance mechanisms. Its controls can be applied selectively across organizational units, which gives the company a consistent security and management foundation while still allowing different projects or compliance regimes to have different account groupings and controls. See [AWS Control Tower documentation](#).

AWS IAM Identity Center is also the appropriate integration point for workforce federation in this architecture. It can be connected to an external SAML 2.0 identity provider such as AD FS, allowing users to authenticate with the existing corporate identity system while gaining centrally managed access to AWS accounts. See [AWS IAM Identity Center external identity provider documentation](#).

QUESTION NO: 43

A company built an ecommerce website on AWS using a three-tier web architecture. The application is Java-based and composed of an Amazon CloudFront distribution, an Apache web server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database.

Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts. The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics. Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis.

Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events? (Choose three.)

- A.** Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs.
- B.** Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java.
- C.** Configure the Aurora MySQL DB cluster to stream slow query and error logs to Amazon Kinesis
- D.** Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs.
- E.** Enable and configure AWS CloudTrail to collect and analyze application activity from Amazon EC2 and Aurora.
- F.** Enable Aurora MySQL DB cluster performance benchmarking and publish the stream to AWS X-Ray.

ANSWER: A B D

Explanation:

Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs is correct because Aurora performance metrics alone do not show the SQL statements that caused latency. Publishing database logs to CloudWatch Logs gives the operations team durable, centralized access to slow query and error information for later analysis, even after the traffic spike has ended. AWS documents this capability for Aurora MySQL log exports to CloudWatch Logs at [Publishing Aurora MySQL logs to Amazon CloudWatch Logs](#).

Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java is also correct because it provides end-to-end request tracing across the Java application and its downstream database calls. This helps identify where time is spent during cart operations and exposes latency, errors, and bottlenecks at the application level. AWS describes Java application instrumentation in [AWS X-Ray SDK for Java](#).

Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs is correct because Auto Scaling can terminate instances before local logs are collected. Streaming Apache access and error logs centrally preserves evidence from short-lived instances and supports searching, metrics, alarms, and dashboards during peak events.

QUESTION NO: 44

A video streaming company recently launched a mobile app for video sharing. The app uploads various files to an Amazon S3 bucket in the us-east-1 Region. The files range in size from 1 GB to 10 GB.

Users who access the app from Australia have experienced uploads that take long periods of time. Sometimes the files fail to completely upload for these users. A solutions architect must improve the app's performance for these uploads.

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

- A. Enable S3 Transfer Acceleration on the S3 bucket. Configure the app to use the Transfer Acceleration endpoint for uploads.
- B. Configure an S3 bucket in each Region to receive the uploads. Use S3 Cross-Region Replication to copy the files to the distribution S3 bucket.
- C. Set up Amazon Route 53 with latency-based routing to route the uploads to the nearest S3 bucket Region.
- D. Configure the app to break the video files into chunks. Use a multipart upload to transfer files to Amazon S3.
- E. Modify the app to add random prefixes to the files before uploading.

ANSWER: A D

Explanation:

Enable S3 Transfer Acceleration on the S3 bucket. Configure the app to use the Transfer Acceleration endpoint for uploads is correct because S3 Transfer Acceleration is designed for long-distance transfers to S3. It routes upload traffic through the nearest AWS edge location and then across the AWS global network to the target S3 bucket, which can significantly improve throughput and reliability for users far from the bucket's Region, such as users in Australia uploading to us-east-1. AWS specifically recommends testing and using this feature when clients are geographically distant from the S3 bucket. See [Amazon S3 Transfer Acceleration](#).

Configure the app to break the video files into chunks. Use a multipart upload to transfer files to Amazon S3 is also correct because multipart upload is the best practice for large objects. For 1 GB to 10 GB video files, multipart upload can upload multiple parts in parallel, improve network utilization, and allow failed parts to be retried without restarting the entire upload. This directly addresses both slow uploads and incomplete uploads. AWS recommends multipart upload for large objects and requires it for objects larger than 5 GB. See [Uploading and copying objects using multipart upload](#).

QUESTION NO: 45

A solutions architect is designing a solution to process events. The solution must have the ability to scale in and out based on the number of events that the solution receives. If a processing error occurs, the event must move into a separate queue for review.

Which solution will meet these requirements?

- A. Send event details to an Amazon Simple Notification Service (Amazon SNS) topic. Configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Add an on-failure destination to the function. Set an Amazon Simple Queue Service (Amazon SQS) queue as the target.

B. Publish events to an Amazon Simple Queue Service (Amazon SQS) queue. Create an Amazon EC2 Auto Scaling group. Configure the Auto Scaling group to scale in and out based on the ApproximateAgeOfOldestMessage metric of the queue. Configure the application to write failed messages to a dead-letter queue.

C. Write events to an Amazon DynamoDB table. Configure a DynamoDB stream for the table. Configure the stream to invoke an AWS Lambda function. Configure the Lambda function to process the events.

D. Publish events to an Amazon EventBridge event bus. Create and run an application on an Amazon EC2 instance with an Auto Scaling group that is behind an Application Load Balancer (ALB). Set the ALB as the event bus target. Configure the event bus to retry events. Write messages to a dead-letter queue if the application cannot process the messages.

ANSWER: A

Explanation:

The solution that sends event details to an Amazon Simple Notification Service topic, invokes an AWS Lambda function as a subscriber, and configures an on-failure destination to an Amazon Simple Queue Service queue is correct. Amazon SNS is a managed publish/subscribe service that can fan out events to Lambda subscribers without requiring the architect to manage polling infrastructure. When SNS invokes Lambda, Lambda can automatically scale concurrency in response to the incoming event volume, which satisfies the requirement to scale in and out based on the number of events received. For failed asynchronous invocations, Lambda supports destinations that can route invocation records to another AWS service after retries are exhausted. Configuring an on-failure destination with Amazon SQS provides a separate queue where failed events can be retained for inspection, troubleshooting, or reprocessing. This design is fully managed, event driven, and directly aligns with both the scaling and failure-handling requirements. See the AWS documentation for [Lambda asynchronous invocation](#) and [using Lambda with Amazon SNS](#).

QUESTION NO: 46

A solutions architect has implemented a SAML 2.0 federated identity solution with their company's on-premises identity provider (IdP) to authenticate users' access to the AWS environment. When the solutions architect tests authentication through the federated identity web portal access to the AWS environment is granted. However, when test users attempt to authenticate through the federated identity web portal, they are not able to access the AWS environment.

Which items should the solutions architect check to ensure identity federation is properly configured? (Select THREE)

- A.** The IAM user's permissions policy has allowed the use of SAML federation for that user
- B.** The IAM roles created for the federated users' or federated groups' trust policy have set the SAML provider as the principle.
- C.** Test users are not in the AWSFederatedUsers group in the company's IdP
- D.** The web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider the ARN of the IAM role, and the SAML assertion from IdP
- E.** The on-premises IdP's DNS hostname is reachable from the AWS environment VPCs.
- F.** The company's IdP defines SAML assertions that properly map users or groups in the company to IAM roles with appropriate permissions

ANSWER: B D F

Explanation:

The correct checks are that the IAM roles created for the federated users' or federated groups' trust policy have set the SAML provider as the principle, that the web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider the ARN of the IAM role, and the SAML assertion from IdP, and that the company's IdP defines SAML assertions that properly map users or groups in the company to IAM roles with appropriate permissions. In a SAML federation design, AWS grants access by allowing a trusted SAML provider to assume IAM roles. Those IAM roles must have a trust policy that identifies the SAML identity provider and permits the sts:AssumeRoleWithSAML action. The federation application or sign-in flow must also submit the SAML assertion and the correct role and provider ARNs to AWS STS so AWS can issue temporary security credentials. Finally, the identity provider must include the required SAML

attributes that map authenticated users or groups to the correct AWS IAM roles; otherwise, users might authenticate successfully to the IdP but receive no usable AWS role. AWS documents these requirements in [IAM SAML 2.0 federation](#) and the [AssumeRoleWithSAML API reference](#).

QUESTION NO: 47

A software as a service (SaaS) company provides a media software solution to customers. The solution is hosted on 50 VPCs across various AWS Regions and AWS accounts. One of the VPCs is designated as a management VPC. The compute resources in the VPCs work independently. The company has developed a new feature that requires all 50 VPCs to be able to communicate with each other. The new feature also requires one-way access from each customer's VPC to the company's management VPC. The management VPC hosts a compute resource that validates licenses for the media software solution. The number of VPCs that the company will use to host the solution will continue to increase as the solution grows. Which combination of steps will provide the required VPC connectivity with the LEAST operational overhead? (Choose two.)

- A. Create a transit gateway. Attach all the company's VPCs and relevant subnets to the transit gateway.
- B. Create VPC peering connections between all the company's VPCs.
- C. Create a Network Load Balancer (NLB) that points to the compute resource for license validation. Create an AWS PrivateLink endpoint service that is available to each customer's VPC. Associate the endpoint service with the NLB.
- D. Create a VPN appliance in each customer's VPC. Connect the company's management VPC to each customer's VPC by using AWS Site-to-Site VPN.
- E. Create a VPC peering connection between the company's management VPC and each customer's VPC.

ANSWER: A C

Explanation:

Create a transit gateway. Attach all the company's VPCs and relevant subnets to the transit gateway. is correct because AWS Transit Gateway is designed as a scalable hub for connecting many VPCs across accounts, avoiding the operational burden of managing large numbers of individual VPC-to-VPC connections as the environment grows. In multi-Region designs, transit gateway peering can extend this model between Regions while keeping routing centralized and easier to operate. This aligns with AWS guidance that Transit Gateway simplifies connectivity for many VPCs and accounts. See [AWS Transit Gateway](#).

Create a Network Load Balancer (NLB) that points to the compute resource for license validation. Create an AWS PrivateLink endpoint service that is available to each customer's VPC. Associate the endpoint service with the NLB. is also correct because AWS PrivateLink is a best-practice pattern for SaaS providers that need to expose a service privately to many consumer VPCs. Customers can initiate private connectivity to the license validation service through interface endpoints, while the provider service remains isolated behind the endpoint service and Network Load Balancer. This provides the required one-way service access with low operational overhead and without requiring broad network-level connectivity to the management VPC. See [AWS PrivateLink](#).

QUESTION NO: 48

A company has a legacy application that runs on multiple .NET Framework components. The components share the same Microsoft SQL Server database and

communicate with each other asynchronously by using Microsoft Message Queueing (MSMQ).

The company is starting a migration to containerized .NET Core components and wants to refactor the application to run on AWS. The .NET Core components require complex orchestration. The company must have full control over networking and host configuration. The application 's database model is strongly relational.

Which solution will meet these requirements?

- A. Host the .NET Core components on AWS App Runner. Host the database on Amazon RDS for SQL Server. Use Amazon EventBridge for asynchronous messaging.

B. Host the .NET Core components on Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Host the database on Amazon DynamoD

C. Host the .NET Core components on AWS Elastic Beanstalk. Host the database on Amazon Aurora PostgreSQL Serverless v2. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) for asynchronous messaging.

D. Host the .NET Core components on Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. Host the database on Amazon Aurora MySQL Serverless v2. Use Amazon Simple Queue Service (Amazon SQS) for asynchronous messaging.

ANSWER: D

Explanation:

Host the .NET Core components on Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type, host the database on Amazon Aurora MySQL Serverless v2, and use Amazon Simple Queue Service (Amazon SQS) for asynchronous messaging is the best fit. Amazon ECS supports sophisticated container orchestration for microservices and batch-style workloads, while the Amazon EC2 launch type gives the company direct control over the underlying container instances. That control is important when teams need to customize host configuration, manage instance-level settings, choose AMIs, or apply specific networking and security configurations beyond what serverless container hosting typically exposes. See the ECS launch type documentation for the distinction between EC2 and managed capacity models: [Amazon ECS launch types](#).

Amazon Aurora MySQL Serverless v2 provides a managed relational database engine that can support a strongly relational application model while scaling capacity automatically. Amazon SQS is also a natural managed replacement for asynchronous queue-based communication patterns such as MSMQ, providing durable message queues that decouple application components. AWS documents SQS as a fully managed message queuing service for decoupling distributed systems and microservices: [Amazon SQS Developer Guide](#).

QUESTION NO: 49

A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalog page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.

Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times.

Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Select TWO.)

A. Configure read replicas for Amazon RDS MySQL and use the single reader endpoint in the web application to reduce the load on the backend database tier.

B. Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionality. Configure Amazon CloudWatch alarms to notify administrators when the site fails.

C. Configure the target group health check to use a TCP check of the Amazon EC2 web server and the Amazon Route S3 health check against the product page to evaluate full application functionality. Configure Amazon CloudWatch alarms to notify administrators when the site fails.

D. Configure an Amazon CloudWatch alarm for Amazon RDS with an action to recover a high-load, impaired RDS instance in the database tier.

E. Configure an Amazon Elastic cache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

ANSWER: B E

Explanation:

Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionality. Configure Amazon CloudWatch alarms to notify administrators when the site fails is correct because it separates infrastructure health from full application dependency health. The ALB target group health check should validate that the EC2 instance and web server can respond, so Auto Scaling does not unnecessarily terminate healthy instances when a downstream database dependency is slow. A separate Route 53 health check against the product page provides end-to-end visibility into the customer-facing application path, and CloudWatch alarms can notify operators when the functional check fails. AWS documents ALB target health checks and Route 53 health checks as mechanisms for validating target and endpoint availability: [Elastic Load Balancing target group health checks](#) and [Route 53 health checks and DNS failover](#).

Configure an Amazon Elastic cache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier is also correct. Caching frequently accessed catalog data reduces repeated reads against MySQL, lowers query latency, and improves scalability for future traffic growth.

QUESTION NO: 50

A software as a service (SaaS) company has developed a multi-tenant environment. The company uses Amazon DynamoDB tables that the tenants share for the storage layer. The company uses AWS Lambda functions for the application services. The company wants to offer a tiered subscription model that is based on resource consumption by each tenant. Each tenant is identified by a unique tenant ID that is sent as part of each request to the Lambda functions. The company has created an AWS Cost and Usage Report (AWS CUR) in an AWS account. The company wants to allocate the DynamoDB costs to each tenant to match that tenant's resource consumption. Which solution will provide a granular view of the DynamoDB cost for each tenant with the LEAST operational effort?

- A.** Associate a new tag that is named tenant ID with each table in DynamoDB. Activate the tag as a cost allocation tag in the AWS Billing and Cost Management console. Deploy new Lambda function code to log the tenant ID in Amazon CloudWatch Logs. Use the AWS CUR to separate DynamoDB consumption cost for each tenant ID.
- B.** Configure the Lambda functions to log the tenant ID and the number of RCUs and WCUs consumed from DynamoDB for each transaction to Amazon CloudWatch Logs. Deploy another Lambda function to calculate the tenant costs by using the logged capacity units and the overall DynamoDB cost from the AWS Cost Explorer API. Create an Amazon EventBridge rule to invoke the calculation Lambda function on a schedule.
- C.** Create a new partition key that associates DynamoDB items with individual tenants. Deploy a Lambda function to populate the new column as part of each transaction. Deploy another Lambda function to calculate the tenant costs by using Amazon Athena to calculate the number of tenant items from DynamoDB and the overall DynamoDB cost from the AWS CUR. Create an Amazon EventBridge rule to invoke the calculation Lambda function on a schedule.
- D.** Deploy a Lambda function to log the tenant ID, the size of each response, and the duration of the transaction call as custom metrics to Amazon CloudWatch Logs. Use CloudWatch Logs Insights to query the custom metrics for each tenant. Use AWS Pricing Calculator to obtain the overall DynamoDB costs and to calculate the tenant costs.
- E.** Configure the Lambda functions to publish per-tenant DynamoDB usage data to AWS Application Cost Profiler. Use Application Cost Profiler reports, backed by AWS cost data, to allocate DynamoDB costs to each tenant.

ANSWER: E

Explanation:

Using AWS Application Cost Profiler is the best fit because it is purpose-built for SaaS environments where multiple tenants consume shared AWS resources. The Lambda functions can capture tenant context and DynamoDB usage for each request, such as the tenant ID, the DynamoDB table resource identifier, and the consumed read and write capacity units. DynamoDB APIs can return consumed capacity information when requests use the ReturnConsumedCapacity parameter, which gives the application accurate usage data to associate with each tenant. Application Cost Profiler then uses the tenant usage data together with AWS cost data to generate tenant-level cost reports, avoiding the need to build and operate a custom cost allocation pipeline. This approach provides the granular per-tenant cost visibility needed for a consumption-based subscription model while minimizing operational overhead. For more details, see [AWS Application Cost Profiler](#) and the DynamoDB documentation for [read/write capacity and consumed capacity](#).

QUESTION NO: 51

A company is planning to migrate its on-premises transaction-processing application to AWS. The application runs inside Docker containers that are hosted on VMs in the company's data center. The Docker containers have shared storage where the application records transaction data. The transactions are time sensitive. The volume of transactions inside the application is unpredictable. The company must implement a low-latency storage solution that will automatically scale throughput to meet increased demand. The company cannot develop the application further and cannot continue to administer the Docker hosting environment. How should the company migrate the application to AWS to meet these requirements?

- A.** Migrate the containers that run the application to Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon S3 to store the transaction data that the containers share.
- B.** Migrate the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS). Create an Amazon Elastic File System (Amazon EFS) file system. Create a Fargate task definition. Add a volume to the task definition to point to the EFS file system.
- C.** Migrate the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS). Create an Amazon Elastic Block Store (Amazon EBS) volume. Create a Fargate task definition. Attach the EBS volume to each running task.
- D.** Launch Amazon EC2 instances. Install Docker on the EC2 instances. Migrate the containers to the EC2 instances. Create an Amazon Elastic File System (Amazon EFS) file system. Add a mount point to the EC2 instances for the EFS file system.

ANSWER: B

Explanation:

Migrate the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS). Create an Amazon Elastic File System (Amazon EFS) file system. Create a Fargate task definition. Add a volume to the task definition to point to the EFS file system. is correct because it combines serverless container operations with a shared, POSIX-compatible file system that can be mounted by multiple running tasks. AWS Fargate removes the need to manage the underlying Docker hosts, capacity, patching, and container infrastructure, which directly satisfies the requirement that the company cannot continue administering the Docker hosting environment. Amazon EFS integrates with Amazon ECS task definitions, allowing containers running on Fargate to mount an EFS file system as a persistent shared volume without application code changes. This is well suited for an existing containerized application that already expects shared file storage. EFS is designed to provide low-latency file access and automatically scales storage capacity as files are added or removed. With EFS throughput modes, including Elastic Throughput, performance can scale with changing workload demand, which fits the unpredictable transaction volume requirement. See the AWS documentation for [Amazon EFS volumes with Amazon ECS](#) and [Amazon EFS performance](#).

QUESTION NO: 52

A global ecommerce company has many data centers worldwide. The company needs scalable cloud storage for legacy file applications. Requirements:

Must support iSCSI access from on-premises servers.

Must support point-in-time snapshots via AWS Backup.

Must retain low-latency access to frequently accessed data. Which solution will meet these requirements?

- A.** Provision an AWS Storage Gateway tape gateway with S3 and AWS Backup.
- B.** Use Amazon FSx File Gateway and S3 File Gateway. Use AWS Backup.
- C.** Provision an AWS Storage Gateway volume gateway in cache mode. Back up the volumes using AWS Backup.
- D.** Provision an AWS Storage Gateway file gateway in cache mode. Use AWS Backup.

ANSWER: C

Explanation:

Provision an AWS Storage Gateway volume gateway in cache mode. Back up the volumes using AWS Backup is correct because Volume Gateway presents cloud-backed block storage to on-premises applications by using the iSCSI protocol, which matches the requirement for legacy servers that need iSCSI access. In cached volume mode, the primary data is stored in Amazon S3 while frequently accessed data is retained locally on the gateway cache. This design provides scalable AWS-backed storage while preserving low-latency access for active working sets in the data center. Volume Gateway also supports point-in-time recovery through snapshots, and AWS Backup can centrally manage backups of Storage Gateway volumes, helping automate retention and backup policies across locations. This combination directly satisfies the required iSCSI connectivity, local low-latency caching, and AWS-managed point-in-time backup capability. For more detail, see the AWS documentation for [Volume Gateway](#) and [backing up Storage Gateway volumes with AWS Backup](#).

QUESTION NO: 53

An ecommerce company runs its infrastructure on AWS. The company exposes its APIs to its web and mobile clients through an Application Load Balancer (ALB) in front of an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster runs thousands of pods that provide the APIs.

After extending delivery to a new continent, the company adds an Amazon CloudFront distribution and sets the ALB as the origin. The company also adds AWS WAF to its architecture.

After implementation of the new architecture, API calls are significantly. However, there is a sudden increase in HTTP status code 504 (Gateway Timeout) errors and HTTP status code 502 (Bad Gateway) errors. This increase in errors seems to be for a specific domain. Which factors could be a cause of these errors? (Select TWO.)

- A. AWS WAF is blocking suspicious requests.
- B. The origin is not properly configured in CloudFront.
- C. There is an SSL/TLS handshake issue between CloudFront and the origin.
- D. EKS Kubernetes pods are being cycled.
- E. Some pods are taking more than 30 seconds to answer API calls.

ANSWER: C E

Explanation:

There is an SSL/TLS handshake issue between CloudFront and the origin is correct because CloudFront returns HTTP 502 Bad Gateway when it cannot negotiate TLS with a custom origin. This can happen if the origin certificate is expired, self-signed, uses an unsupported protocol or cipher, or the certificate name does not match the domain name that CloudFront uses to connect to the origin. This aligns especially well with errors that appear for a specific domain. See the AWS guidance for troubleshooting CloudFront 502 errors at [AWS CloudFront 502 Bad Gateway](#).

Some pods are taking more than 30 seconds to answer API calls is also correct because CloudFront has an origin response timeout, which is 30 seconds by default. If the ALB or backend application pods do not return a response before that timeout expires, CloudFront can return HTTP 504 Gateway Timeout to the viewer. AWS documents that 504 errors occur when the origin does not respond before the request expires, including latency or application response delays behind the origin. See [AWS CloudFront 504 Gateway Timeout](#).

QUESTION NO: 54

A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS. Which solution will meet these requirements?

- A. Migrate the SQL Server databases to Amazon RDS for MySQL by using backup and restore utilities.
- B. Use an AWS Snowball Edge Storage Optimized device to transfer data to Amazon S3. Set up Amazon RDS for MySQL. Use S3 integration with SQL Server features, such as BULK INSERT.

C. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MySQL. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.

D. Use AWS DataSync to migrate data over the network between on-premises storage and Amazon S3. Set up Amazon RDS for MySQL. Use S3 integration with SQL Server features, such as BULK INSERT.

ANSWER: C

Explanation:

Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MySQL. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS is correct because the source and target database engines are different. Moving from Microsoft SQL Server to MySQL is a heterogeneous database migration, which requires schema and code conversion in addition to moving table data. AWS Schema Conversion Tool is designed for this purpose: it assesses and converts source database schemas, views, stored procedures, and other database objects into a format compatible with the target engine where possible. After the target schema is prepared in Amazon RDS for MySQL, AWS Database Migration Service can migrate the data from the on-premises SQL Server environment to the AWS managed database with minimal downtime, and it can support ongoing replication during cutover. This pairing is the standard AWS approach for heterogeneous migrations to managed relational databases such as Amazon RDS. For more details, see the AWS documentation for [AWS Schema Conversion Tool](#) and [AWS Database Migration Service](#).

QUESTION NO: 55

A company needs to build a disaster recovery (DR) solution for its ecommerce website. The web application is hosted on a fleet of t3.large Amazon EC2 instances and uses an Amazon RDS for MySQL DB instance. The EC2 instances are in an Auto Scaling group that extends across multiple Availability Zones.

In the event of a disaster, the web application must fail over to the secondary environment with an RPO of 30 seconds and an RTO of 10 minutes.

Which solution will meet these requirements MOST cost-effectively?

A. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create a cross-Region read replica for the DB instance. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region. Recover the EC2 instances from the latest EC2 backup. Use an Amazon Route 53 geoloc

B. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create a cross-Region read replica for the DB instance. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. Run the EC2 instances at the minimum capacity in the DR Region Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. Increase the desired

C. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Manually restore the backed-up data on new instances. Use an Amazon Route 53 simple routing policy to automatically fail over to the DR Reg

D. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create an Amazon Aurora global database. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. Run the Auto Scaling group of EC2 instances at full capacity in the DR Region. Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster.

ANSWER: B

Explanation:

The solution that uses infrastructure as code, an Amazon RDS cross-Region read replica, AWS Elastic Disaster Recovery, minimum EC2 capacity in the DR Region, and Amazon Route 53 failover routing is correct because it provides a warm standby-style architecture that balances recovery speed and cost. AWS Elastic Disaster Recovery continuously replicates source EC2 instance block storage to a staging area in the target Region and can launch recovery instances within minutes, which aligns with the 10-minute RTO requirement. Keeping only minimum EC2 capacity running in the DR Region avoids the

expense of maintaining a full duplicate production environment while still allowing the Auto Scaling group capacity to be increased during failover. For the database tier, an RDS for MySQL cross-Region read replica provides ongoing asynchronous replication to the DR Region, supporting a low RPO target for regional disaster recovery. Route 53 failover routing can direct users to the DR environment when health checks detect that the primary environment is unavailable. This combination is a cost-effective fit for a 30-second RPO and 10-minute RTO. References: [AWS Elastic Disaster Recovery](#) and [Amazon Route 53 failover routing](#).

QUESTION NO: 56

A solution architect is designing an AWS account structure for a company that consists of multiple terms. All the team will work in the same AWS Region. The company needs a VPC that is connected to the on-premises network. The company expects less than 50 Mbps of total to and from the on-premises network.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO)

- A. Create an AWS CloudFormation template that provisions a VPC and the required subnets. Deploy the template to each AWS account
- B. Create an AWS CloudFormation template that provisions a VPC and the required subnets. Deploy the template to a shared services account. Share the subnets by using AWS Resource Access Manager
- C. Use AWS Transit Gateway along with an AWS Site-to-Site VPN for connectivity to the on-premises network. Share the transit gateway by using AWS Resource Access Manager
- D. Use AWS Site-to-Site VPN for connectivity to the on-premises network
- E. Use AWS Direct Connect for connectivity to the on-premises network.

ANSWER: B D

Explanation:

Create an AWS CloudFormation template that provisions a VPC and the required subnets. Deploy the template to a shared services account. Share the subnets by using AWS Resource Access Manager is correct because a shared VPC model lets multiple AWS accounts deploy resources into centrally managed subnets without duplicating VPCs, route tables, gateways, and connectivity in every account. AWS Resource Access Manager supports VPC subnet sharing within an AWS Organization, which is well suited when multiple teams operate in the same Region and need access to a common network foundation. This reduces operational overhead and avoids unnecessary per-account network duplication. See [AWS VPC sharing](#).

Use AWS Site-to-Site VPN for connectivity to the on-premises network is also correct because the expected traffic is less than 50 Mbps, making an internet-based VPN connection a cost-effective way to connect the shared VPC to the on-premises environment. Site-to-Site VPN provides encrypted IPsec tunnels and can be attached to a virtual private gateway for VPC connectivity, which fits a modest bandwidth requirement without the higher cost and provisioning effort of a dedicated private connection. See [AWS Site-to-Site VPN](#).

QUESTION NO: 57

A company manufactures smart vehicles. The company uses a custom application to collect vehicle data. The vehicles use the MQTT protocol to connect to the application. The company processes the data in 5-minute intervals. The company then copies vehicle telematics data to on-premises storage. Custom applications analyze this data to detect anomalies. The number of vehicles that send data grows constantly. Newer vehicles generate high volumes of data. The on-premises storage solution is not able to scale for peak traffic, which results in data loss. The company must modernize the solution and migrate the solution to AWS to resolve the scaling challenges. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS IoT Greengrass to send the vehicle data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create an Apache Kafka application to store the data in Amazon S3. Use a pretrained model in Amazon SageMaker to detect anomalies.

B. Use AWS IoT Core to receive the vehicle data. Configure rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3. Create an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies.

C. Use AWS IoT FleetWise to collect the vehicle data. Send the data to an Amazon Kinesis data stream. Use an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use the built-in machine learning transforms in AWS Glue to detect anomalies.

D. Use Amazon MQ for RabbitMQ to collect the vehicle data. Send the data to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Lookout for Metrics to detect anomalies.

ANSWER: B

Explanation:

Use AWS IoT Core to receive the vehicle data. Configure rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3. Create an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies. is correct because it uses fully managed AWS services that directly match the workload requirements. AWS IoT Core natively supports MQTT, which allows the vehicles to connect securely without the company operating broker infrastructure. IoT rules can route incoming messages directly to AWS services, including Kinesis Data Firehose, reducing custom integration work and operational overhead. Kinesis Data Firehose can buffer, batch, scale automatically, and deliver high-volume streaming telemetry data into Amazon S3, which provides durable, elastic storage for downstream analytics. For the 5-minute processing and anomaly detection requirement, Amazon Kinesis Data Analytics can process streaming data and includes SQL-based analytics capabilities such as anomaly detection using Random Cut Forest, allowing the company to analyze data in near real time without managing servers. This architecture replaces the non-scaling on-premises storage path with managed ingestion, streaming, storage, and analytics components. Relevant AWS documentation includes [AWS IoT rules](#) and [Kinesis Data Analytics anomaly detection](#).

QUESTION NO: 58

A company has more than 10,000 sensors that send data to an on-premises Apache Kafka server by using the Message Queuing Telemetry Transport (MQTT) protocol. The on-premises Kafka server transforms the data and then stores the results as objects in an Amazon S3 bucket.

Recently, the Kafka server crashed. The company lost sensor data while the server was being restored. A solutions architect must create a new design on AWS that is highly available and scalable to prevent a similar occurrence.

Which solution will meet these requirements?

A. Launch two Amazon EC2 instances to host the Kafka server in an active/standby configuration across two Availability Zones. Create a domain name in Amazon Route 53. Create a Route 53 failover policy. Route the sensors to send the data to the domain name.

B. Migrate the on-premises Kafka server to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create a Network Load Balancer (NLB) that points to the Amazon MSK broker. Enable NLB health checks. Route the sensors to send the data to the NLB.

C. Deploy AWS IoT Core, and connect it to an Amazon Kinesis Data Firehose delivery stream. Use an AWS Lambda function to handle data transformation. Route the sensors to send the data to AWS IoT Core.

D. Deploy AWS IoT Core, and launch an Amazon EC2 instance to host the Kafka server. Configure AWS IoT Core to send the data to the EC2 instance. Route the sensors to send the data to AWS IoT Core.

ANSWER: C

Explanation:

Deploy AWS IoT Core, and connect it to an Amazon Kinesis Data Firehose delivery stream. Use an AWS Lambda function to handle data transformation. Route the sensors to send the data to AWS IoT Core is correct because it replaces the single on-premises ingestion and transformation point with fully managed, horizontally scalable AWS services. AWS IoT Core is designed for device connectivity and supports MQTT, making it a natural fit for more than 10,000 sensors that already use the MQTT protocol. It can securely receive device messages and use IoT rules to route those messages to downstream

AWS services without the company managing broker infrastructure. Amazon Kinesis Data Firehose is also fully managed and can deliver streaming data to Amazon S3 while buffering, scaling, retrying delivery, and optionally invoking AWS Lambda for record transformation before the data is written. This architecture removes the Kafka server as a single point of failure and provides durable, highly available ingestion and delivery to S3 with minimal operational overhead. For reference, see [AWS IoT Core documentation](#) and [Amazon Data Firehose documentation](#).

QUESTION NO: 59

A company uses AWS Organizations to manage more than 1,000 AWS accounts. The company has created a new developer organization. There are 540 developer member accounts that must be moved to the new developer organization. All accounts are set up with all the required information so that each account can be operated as a standalone account. Which combination of steps should a solutions architect take to move all of the developer accounts to the new developer organization? (Choose three.)

- A. Call the MoveAccount operation in the Organizations API from the old organization's management account to migrate the developer accounts to the new developer organization.
- B. From the management account, remove each developer account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.
- C. From each developer account, remove the account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.
- D. Sign in to the new developer organization's management account and create a placeholder member account that acts as a target for the developer account migration.
- E. Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization's management account to send invitations to the developer accounts.
- F. Have each developer sign in to their account and confirm to join the new developer organization.

ANSWER: B E F

Explanation:

The correct process is to make each developer account standalone first, then invite the standalone accounts into the new AWS Organization, and finally have the accounts accept the invitations. "From the management account, remove each developer account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API" is correct because AWS Organizations allows the management account to remove member accounts, provided each account has the required standalone account information such as contact details, a payment method, and agreement acceptance. "Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization's management account to send invitations to the developer accounts" is correct because the new organization's management account must initiate invitations for existing standalone AWS accounts to join. "Have each developer sign in to their account and confirm to join the new developer organization" is also correct because invited accounts must accept the handshake before becoming members of the new organization. This aligns with the documented AWS Organizations workflow for removing member accounts and inviting AWS accounts into an organization. See the AWS documentation for [removing member accounts from an organization](#) and [inviting AWS accounts to join an organization](#).

QUESTION NO: 60

A solutions architect is importing a VM from an on-premises environment by using the Amazon EC2 VM Import feature of AWS Import/Export. The solutions architect has created an AMI and has provisioned an Amazon EC2 instance that is based on that AMI. The EC2 instance runs inside a public subnet in a VPC and has a public IP address assigned.

The EC2 instance does not appear as a managed instance in the AWS Systems Manager console.

Which combination of steps should the solutions architect take to troubleshoot this issue? (Select TWO.)

- A. Verify that Systems Manager Agent is installed on the instance and is running.
- B. Verify that the instance is assigned an appropriate IAM role for Systems Manager.

- C. Verify the existence of a VPC endpoint on the VPC.
- D. Verify that the AWS Application Discovery Agent is configured.
- E. Verify the correct configuration of service-linked roles for Systems Manager.

ANSWER: A B

Explanation:

Verify that Systems Manager Agent is installed on the instance and is running is correct because an EC2 instance must run the SSM Agent before it can register with AWS Systems Manager as a managed node. This is especially important for an imported VM, because the source on-premises operating system image might not include the agent, might include an outdated version, or might not have the agent service enabled. AWS documents SSM Agent as a core prerequisite for using Systems Manager with EC2 instances: [Working with SSM Agent](#).

Verify that the instance is assigned an appropriate IAM role for Systems Manager is also correct because the agent must authenticate to Systems Manager and call required AWS service APIs. For EC2 instances, AWS recommends attaching an instance profile with a role that includes the AmazonSSMManagedInstanceCore managed policy. That policy grants the permissions needed for Systems Manager core functionality, including managed node registration and communication with Systems Manager service endpoints. AWS lists the instance profile permissions as a required setup step for managed EC2 instances: [Configure instance permissions for Systems Manager](#).

QUESTION NO: 61

A company has used infrastructure as code (IaC) to provision a set of two Amazon EC2 instances. The instances have remained the same for several years.

The company's business has grown rapidly in the past few months. In response, the company's operations team has implemented an Auto Scaling group to manage the sudden increases in traffic. Company policy requires a monthly installation of security updates on all operating systems that are running.

The most recent security update required a reboot. As a result the Auto Scaling group terminated the instances and replaced them with new, unpatched instances.

Which combination of steps should a solutions architect recommend to avoid a recurrence of this issue? (Select TWO)

- A. Modify the Auto Scaling group by setting the Update policy to target the oldest launch configuration for replacement.
- B. Create a new Auto Scaling group before the next patch maintenance. During the maintenance window, patch both groups and reboot the instances.
- C. Create an Elastic Load Balancer in front of the Auto Scaling group. Configure monitoring to ensure that target group health checks return healthy after the Auto Scaling group replaces the terminated instances.
- D. Create automation scripts to patch an AMI, update the launch configuration, and invoke an Auto Scaling instance refresh.
- E. Create an Elastic Load Balancer in front of the Auto Scaling group. Configure termination protection on the instances.

ANSWER: A D

Explanation:

Create automation scripts to patch an AMI, update the launch configuration, and invoke an Auto Scaling instance refresh. is correct because the root cause is that replacement instances are being launched from an outdated image or configuration. The scalable, repeatable fix is to bake the required security updates into a new AMI, update the Auto Scaling group's launch configuration or launch template to reference that AMI, and then use an instance refresh to roll the fleet forward in a controlled way. Amazon EC2 Auto Scaling instance refresh is designed for this type of rolling replacement when a launch template or launch configuration changes. See [AWS EC2 Auto Scaling instance refresh](#).

Modify the Auto Scaling group by setting the Update policy to target the oldest launch configuration for replacement. is also correct because targeting instances that were launched from the oldest launch configuration helps phase out instances

based on the obsolete, unpatched configuration first. This aligns the Auto Scaling group's replacement behavior with the goal of eliminating older instances and ensuring that future capacity comes from the updated, patched configuration. AWS documents the OldestLaunchConfiguration termination policy as useful when updating a group and replacing instances that use a previous launch configuration. See [Amazon EC2 Auto Scaling termination policies](#).

QUESTION NO: 62

A solutions architect is designing a network for a new cloud deployment. Each account will need autonomy to modify route tables and make changes. Centralized and controlled egress internet connectivity is also needed. The cloud footprint is expected to grow to thousands of AWS accounts.

Which architecture will meet these requirements?

- A.** A centralized transit VPC with a VPN connection to a standalone VPC in each account. Outbound internet traffic will be controlled by firewall appliances.
- B.** A centralized shared VPC with a subnet for each account. Outbound internet traffic will be controlled through a fleet of proxy servers.
- C.** A shared services VPC to host central assets to include a fleet of firewalls with a route to the internet. Each spoke VPC will peer to the central VPC.
- D.** A shared transit gateway to which each VPC will be attached. Outbound internet access will route through a fleet of VPN-attached firewalls.

ANSWER: D

Explanation:

A shared transit gateway to which each VPC will be attached. Outbound internet access will route through a fleet of VPN-attached firewalls is correct because AWS Transit Gateway is designed as a scalable hub-and-spoke routing service for large multi-account environments. It can be shared across accounts by using AWS Resource Access Manager, allowing individual account owners to attach their VPCs and manage the VPC route tables that send default traffic toward the transit gateway. This preserves account-level autonomy while still enabling central network teams to control routing domains through transit gateway route tables.

For centralized and controlled internet egress, spoke VPCs can route outbound traffic to the transit gateway, which then forwards traffic to a centralized inspection or egress environment containing firewall appliances and internet connectivity. This is the standard AWS pattern for scaling outbound inspection and egress control across many VPCs and accounts. AWS explicitly positions Transit Gateway as a way to simplify connectivity at scale and supports centralized egress architectures for multi-VPC deployments. See the AWS documentation for [AWS Transit Gateway](#) and the AWS whitepaper section on [centralized egress to the internet](#).

QUESTION NO: 63

A company's AWS environment includes an Amazon RDS for MySQL database in a Multi-AZ deployment and an Amazon EC2 Auto Scaling group behind an Application Load Balancer (ALB). The Auto Scaling group spans two Availability Zones. The company also uses Amazon Route 53 for DNS hosting.

The company runs an application in its AWS environment. More than 95% of the application's operations are read operations. A solutions architect needs to deploy the workload to a second AWS Region. The solution must reduce application latency while maintaining business continuity.

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

- A.** Migrate the RDS for MySQL database to an Amazon Aurora MySQL global database. Create an ALB in the new Region. Deploy a new EC2 Auto Scaling group behind the new ALB.
- B.** Migrate the RDS for MySQL database to a Multi-AZ deployment in a new Region. Create an ALB in the new Region. Deploy an Amazon CloudFront distribution in front of the new ALB.
- C.** Configure latency-based routing in Route 53. Add a new record that points to both ALBs.

D. Configure geolocation routing in Route 53. Add a new alias record that points to both ALBs.

E. Migrate the RDS for MySQL database to Amazon Aurora Serverless v2. Create a new ALB. Deploy an EC2 Auto Scaling group behind the new ALB.

ANSWER: A C

Explanation:

The correct combination is to migrate the database to an Amazon Aurora MySQL global database and deploy the application stack in the second Region, then use Route 53 latency-based routing to direct users to the lowest-latency Application Load Balancer. Aurora Global Database is designed for cross-Region architectures with one primary Region and secondary Regions that can serve low-latency read traffic. Because more than 95% of the application operations are reads, placing Aurora read capability closer to users in another Region directly reduces database access latency while preserving a managed replication model and supporting disaster recovery through secondary Region promotion. AWS documents Aurora Global Database as supporting globally distributed applications with low-latency global reads and cross-Region disaster recovery: [Amazon Aurora Global Database](#).

Route 53 latency-based routing complements this design by routing DNS queries to the Region that provides the best latency for the user. Pointing latency records at the ALBs in both Regions lets the application tier serve users from the closest healthy regional endpoint, helping reduce end-user response time while maintaining business continuity. AWS describes this routing policy here: [Latency-based routing](#).

QUESTION NO: 64

A company is migrating to the cloud. It wants to evaluate the configurations of virtual machines in its existing data center environment to ensure that it can size new Amazon EC2 instances accurately. The company wants to collect metrics, such as CPU, memory, and disk utilization, and it needs an inventory of what processes are running on each instance. The company would also like to monitor network connections to map communications between servers. Which would enable the collection of this data MOST cost effectively?

- A. Use AWS Application Discovery Service and deploy the data collection agent to each virtual machine in the data center.
- B. Configure the Amazon CloudWatch agent on all servers within the local environment and publish metrics to Amazon CloudWatch Logs.
- C. Use AWS Application Discovery Service and enable agentless discovery in the existing virtualization environment.
- D. Enable AWS Application Discovery Service in the AWS Management Console and configure the corporate firewall to allow scans over a VPN.

ANSWER: A

Explanation:

Use AWS Application Discovery Service and deploy the data collection agent to each virtual machine in the data center is correct because the AWS Application Discovery Agent is designed for detailed, host-level discovery during migration planning. When installed on each server, the agent collects system configuration, performance utilization such as CPU, memory, and disk usage, running processes, and inbound/outbound network connection information. This data helps build dependency maps between servers and supports more accurate right-sizing of target Amazon EC2 instances. It is also cost effective because AWS Application Discovery Service does not charge for collecting and storing discovery data; customers generally only pay for related AWS resources they choose to use, such as migration tooling or storage beyond the service's included capabilities. This makes the agent-based Application Discovery Service approach a strong fit when detailed process and network dependency data is required for migration assessment. See the AWS documentation for the [AWS Application Discovery Agent](#) and [AWS Application Discovery Service pricing](#).

QUESTION NO: 65

A company has set up its entire infrastructure on AWS. The company uses Amazon EC2 instances to host its ecommerce website and uses Amazon S3 to store static data. Three engineers at the company handle the cloud administration and development through one AWS account. Occasionally, an engineer alters an EC2 security group configuration of another engineer and causes noncompliance issues in the environment. A solutions architect must set up a system that tracks changes that the engineers make. The system must send alerts when the engineers make noncompliant changes to the security settings for the EC2 instances. What is the FASTEST way for the solutions architect to meet these requirements?

- A.** Set up AWS Organizations for the company. Apply SCPs to govern and track noncompliant security group changes that are made to the AWS account.
- B.** Enable AWS CloudTrail to capture the changes to EC2 security groups. Enable Amazon CloudWatch rules to provide alerts when noncompliant security settings are detected.
- C.** Enable SCPs on the AWS account to provide alerts when noncompliant security group changes are made to the environment.
- D.** Enable AWS Config on the EC2 security groups to track any noncompliant changes. Send the changes as alerts through an Amazon Simple Notification Service (Amazon SNS) topic.

ANSWER: D

Explanation:

Enable AWS Config on the EC2 security groups to track any noncompliant changes. Send the changes as alerts through an Amazon Simple Notification Service (Amazon SNS) topic is the correct choice because AWS Config is purpose-built to continuously record configuration changes and evaluate resources against compliance rules. For EC2 security groups, AWS Config can capture configuration history, show who or what changed a resource through its integration with AWS CloudTrail data, and evaluate the security group against managed or custom AWS Config rules. When a rule finds that a security group violates the desired security posture, AWS Config can publish notifications through Amazon SNS, allowing the company to alert administrators quickly without building a custom detection workflow. This is the fastest operational approach because it uses native compliance evaluation and notification capabilities rather than only collecting raw API events. AWS documentation describes AWS Config as a service that assesses, audits, and evaluates configurations of AWS resources, and it supports notifications for configuration changes and compliance state changes. See [What is AWS Config?](#) and [Notifications that AWS Config sends to an Amazon SNS topic](#).

QUESTION NO: 66

A company has multiple AWS accounts. The company recently had a security audit that revealed many unencrypted Amazon Elastic Block Store (Amazon EBS) volumes attached to Amazon EC2 instances.

A solutions architect must encrypt the unencrypted volumes and ensure that unencrypted volumes will be detected automatically in the future. Additionally, the company wants a solution that can centrally manage multiple AWS accounts with a focus on compliance and security.

Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

- A.** Create an organization in AWS Organizations. Set up AWS Control Tower, and turn on the strongly recommended guardrails. Join all accounts to the organization. Categorize the AWS accounts into OUs.
- B.** Use the AWS CLI to list all the unencrypted volumes in all the AWS accounts. Run a script to encrypt all the unencrypted volumes in place.
- C.** Create a snapshot of each unencrypted volume. Create a new encrypted volume from the unencrypted snapshot. Detach the existing volume, and replace it with the encrypted volume.
- D.** Create an organization in AWS Organizations. Set up AWS Control Tower, and turn on the mandatory guardrails. Join all accounts to the organization. Categorize the AWS accounts into OUs.
- E.** Turn on AWS CloudTrail. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to detect and automatically encrypt unencrypted volumes.

ANSWER: A C

Explanation:

The correct approach is to use AWS Organizations with AWS Control Tower strongly recommended guardrails for centralized multi-account governance, and to replace each existing unencrypted EBS volume with an encrypted volume created from a snapshot. AWS Control Tower is designed to establish and govern a secure, compliant multi-account AWS environment. Its strongly recommended controls include detective controls for security best practices, such as detecting whether Amazon EBS volumes attached to Amazon EC2 instances are encrypted, which directly addresses the requirement to automatically identify future noncompliant resources across accounts and organizational units. See the AWS Control Tower controls reference at [AWS Control Tower controls](#).

For existing unencrypted EBS volumes, encryption cannot simply be enabled in place on the attached volume. The standard remediation pattern is to create a snapshot of the unencrypted volume, create an encrypted volume from that snapshot, detach the original volume during an appropriate maintenance window, and attach the encrypted replacement volume. Amazon EBS supports creating encrypted volumes from snapshots, allowing the company to remediate the audited resources while preserving the data. See [Amazon EBS encryption](#) for details.

QUESTION NO: 67

A company needs to gather data from an experiment in a remote location that does not have internet connectivity. During the experiment, sensors that are connected to a local network will generate 6 TB of data in a preprimary format over the course of 1 week. The sensors can be configured to upload their data files to an FTP server periodically, but the sensors do not have their own FTP server. The sensors also do not support other protocols. The company needs to collect the data centrally and move the data to object storage in the AWS Cloud as soon as possible after the experiment.

Which solution will meet these requirements?

- A.** Order an AWS Snowball Edge Compute Optimized device. Connect the device to the local network. Configure AWS DataSync with a target bucket name, and unload the data over NFS to the device. After the experiment return the device to AWS so that the data can be loaded into Amazon S3.
- B.** Order an AWS Snowcone device, including an Amazon Linux 2 AMI. Connect the device to the local network. Launch an Amazon EC2 instance on the device. Create a shell script that periodically downloads data from each sensor. After the experiment, return the device to AWS so that the data can be loaded as an Amazon Elastic Block Store [Amazon EBS] volume.
- C.** Order an AWS Snowcone device, including an Amazon Linux 2 AMI. Connect the device to the local network. Launch an Amazon EC2 instance on the device. Install and configure an FTP server on the EC2 instance. Configure the sensors to upload data to the EC2 instance. After the experiment, return the device to AWS so that the data can be loaded into Amazon S3.
- D.** Order an AWS Snowcone device. Connect the device to the local network. Configure the device to use Amazon FSx. Configure the sensors to upload data to the device. Configure AWS DataSync on the device to synchronize the uploaded data with an Amazon S3 bucket. Return the device to AWS so that the data can be loaded as an Amazon Elastic Block Store (Amazon EBS) volume.

ANSWER: C

Explanation:

Order an AWS Snowcone device, including an Amazon Linux 2 AMI. Connect the device to the local network. Launch an Amazon EC2 instance on the device. Install and configure an FTP server on the EC2 instance. Configure the sensors to upload data to the EC2 instance. After the experiment, return the device to AWS so that the data can be loaded into Amazon S3 is correct because it satisfies the disconnected-site constraint, the protocol limitation, and the need to move the collected data into AWS object storage. AWS Snowcone is designed for edge locations and disconnected environments, provides usable local storage capacity sufficient for 6 TB of data, and supports running Amazon EC2-compatible instances locally. By launching an Amazon Linux 2 instance on the Snowcone device and installing an FTP server, the company can provide the exact upload endpoint that the sensors require without needing internet connectivity or a separate server. After the experiment, the Snowcone device can be shipped back to AWS, where the data from an import job can be transferred into Amazon S3. This approach uses Snowcone both as local compute and storage at the edge and as the offline data transfer

mechanism into AWS. See the AWS Snowcone overview in the [AWS Snowcone User Guide](#) and AWS guidance for [using EC2-compatible instances on Snow Family devices](#).

QUESTION NO: 68

A solutions architect is designing a publicly accessible web application that is on an Amazon CloudFront distribution with an Amazon S3 website endpoint as the origin. When the solution is deployed, the website returns an Error 403: Access Denied message.

Which steps should the solutions architect take to correct the issue? (Select TWO.)

- A. Remove the S3 block public access option from the S3 bucket.
- B. Remove the requester pays option from the S3 bucket.
- C. Remove the origin access identity (OAI) from the CloudFront distribution.
- D. Change the storage class from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA).
- E. Disable S3 object versioning.

ANSWER: A B

Explanation:

For a CloudFront distribution that uses an Amazon S3 static website endpoint as the origin, the S3 content must be reachable through anonymous public access. Therefore,

Remove the S3 block public access option from the S3 bucket.

is correct because S3 Block Public Access settings can prevent a bucket policy or ACL from granting public read access to the website objects, resulting in Access Denied responses.

Remove the requester pays option from the S3 bucket.

is also correct because Requester Pays buckets require the requester to include billing-related request information, which is not compatible with anonymous public website access through an S3 website endpoint and can cause 403 errors. AWS guidance for serving static websites through CloudFront distinguishes between using an S3 REST API endpoint with an origin access identity or origin access control, and using an S3 website endpoint with public access allowed. See AWS guidance on [serving static websites through CloudFront](#) and the Amazon S3 documentation for [permissions required for website access](#).

QUESTION NO: 69

A company has 50 AWS accounts that are members of an organization in AWS Organizations. Each account contains multiple VPCs. The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account. Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.

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

- A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager
- B. From the management account, share the transit gateway with member accounts by using an AWS Organizations SCP
- C. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a VPC transit gateway attachment in a member account. Associate the attachment with the transit gateway in the management account by using the transit gateway role
- D. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a peering transit gateway attachment in a member account. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.

E. From the management account, share the transit gateway with member accounts by using AWS Service Catalog

ANSWER: A C

Explanation:

The correct combination is to share the transit gateway from the management account with member accounts by using AWS Resource Access Manager, and to launch an AWS CloudFormation StackSet from the management account that automatically creates the new VPC and VPC transit gateway attachment in each member account. AWS Resource Access Manager is the standard AWS service for sharing a transit gateway across accounts in an AWS Organization, allowing member accounts to create attachments to the shared transit gateway. AWS CloudFormation StackSets is designed for centralized, repeatable deployment of infrastructure across multiple AWS accounts and Regions, and it can automatically deploy stack instances to accounts that are added to an organizational unit. Together, these services support the required operating model: centralized transit gateway ownership, shared network access for member accounts, and automated provisioning of VPC infrastructure and transit gateway attachments when new accounts are created. The VPC transit gateway attachment can be defined with AWS CloudFormation using the appropriate EC2 transit gateway attachment resource. For more information, see [Amazon VPC Transit Gateways](#) and [AWS CloudFormation StackSets](#).

QUESTION NO: 70

A large company with hundreds of AWS accounts has a newly established centralized internal process for purchasing new or modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement or execution. Previously, business units would directly purchase or modify Reserved Instances in their own respective AWS accounts autonomously.

Which combination of steps should be taken to proactively enforce the new process in the MOST secure way possible? (Select TWO.)

- A. Ensure all AWS accounts are part of an AWS Organizations structure operating in all features mode.
- B. Use AWS Config to report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
- C. In each AWS account, create an IAM policy with a DENY rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
- D. Create an SCP that contains a deny rule to the ec2:PurchaseReservedInstancesOffering and ec2: Modify Reserved Instances actions. Attach the SCP to each organizational unit (OU) of the AWS Organizations structure.
- E. Ensure that all AWS accounts are part of an AWS Organizations structure operating in consolidated billing features mode.

ANSWER: A D

Explanation:

The correct combination is to ensure all AWS accounts are part of an AWS Organizations structure operating in all features mode and to create an SCP that contains a deny rule to the Reserved Instance purchase and modification actions, then attach that SCP to the relevant organizational units. AWS Organizations must be enabled with all features to use service control policies, which provide centralized permission guardrails across member accounts. This is the right foundation for a large multi-account environment because it allows the company to enforce the new procurement process consistently instead of relying on each account to configure its own controls.

An SCP that denies the EC2 actions for purchasing Reserved Instance offerings and modifying Reserved Instances proactively prevents business units from bypassing the centralized team. SCPs apply maximum available permissions to accounts and principals within the targeted organization hierarchy, so even highly privileged IAM principals in member accounts cannot perform actions explicitly denied by the SCP, except for identities outside SCP scope such as the management account root user. See the AWS documentation for [service control policies](#) and [enabling all features in AWS Organizations](#).

QUESTION NO: 71

A solutions architect is importing a VM from an on-premises environment by using the Amazon EC2 VM Import feature of AWS Import/Export. The solutions architect has created an AMI and has provisioned an Amazon EC2 instance that is based on that AMI. The EC2 instance runs inside a public subnet in a VPC and has a public IP address assigned. The EC2 instance does not appear as a managed instance in the AWS Systems Manager console. Which combination of steps should the solutions architect take to troubleshoot this issue? (Choose two.)

- A. Verify that Systems Manager Agent is installed on the instance and is running.
- B. Verify that the instance is assigned an appropriate IAM role for Systems Manager.
- C. Verify the existence of a VPC endpoint on the VPC.
- D. Verify that the AWS Application Discovery Agent is configured.
- E. Verify the correct configuration of service-linked roles for Systems Manager.

ANSWER: A B

Explanation:

To appear as a managed instance in AWS Systems Manager, an EC2 instance must meet the core Systems Manager prerequisites. “Verify that Systems Manager Agent is installed on the instance and is running” is correct because imported VMs do not necessarily include a current, functioning SSM Agent. Systems Manager relies on this agent to register the instance, receive commands, and report status back to the service. If the agent is missing, stopped, outdated, or unable to start, the instance will not appear as managed.

“Verify that the instance is assigned an appropriate IAM role for Systems Manager” is also correct because the instance needs permissions to communicate with Systems Manager APIs. AWS recommends attaching an instance profile that includes the AmazonSSMManagedInstanceCore managed policy, which grants the minimum permissions required for Systems Manager core functionality. Since the instance is in a public subnet with a public IP address, it can normally reach the public Systems Manager endpoints through the internet route, assuming normal network controls allow it. The key troubleshooting checks in this scenario are therefore the SSM Agent state and the EC2 instance profile permissions. See the AWS documentation for [Systems Manager managed node prerequisites](#) and [configuring instance permissions for Systems Manager](#).

QUESTION NO: 72

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 1 business day, with loss of no more than 1 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 the CloudFormation 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 secondary Region.

ANSWER: C

Explanation:

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 is correct because AWS Backup provides a managed, policy-based service for protecting Amazon EC2 instances and their attached Amazon EBS volumes. A daily backup schedule satisfies the requirement for no more than one day of data loss, and cross-Region backup copy satisfies the requirement to maintain backups in a separate AWS Region. AWS Backup can back up EC2 instances in a way that captures the instance and associated EBS volume data, and it supports restore operations from recovery points, reducing the amount of manual snapshot orchestration that a small operations team would need to perform. Combining the existing AWS CloudFormation template for the secondary Region network with AWS Backup recovery points provides a practical recovery workflow within one business day while keeping the solution cost-efficient because standby compute resources do not need to run continuously. For details, see [AWS Backup documentation](#) and [AWS Backup cross-Region backup](#).

QUESTION NO: 73

A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.

A solutions architect has created an IAM role that is named `strategy_reviewer` in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Account.

The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

- A.** Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to the account ID of the Strategy account
- B.** Update the `strategy_reviewer` IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.
- C.** Update the custom KMS key policy in the Creative account to grant decrypt permissions to the `strategy_reviewer` IAM role.
- D.** Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to an anonymous user.
- E.** Update the custom KMS key policy in the Creative account to grant encrypt permissions to the `strategy_reviewer` IAM role.
- F.** Update the `strategy_reviewer` IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key

ANSWER: A C F

Explanation:

For cross-account access to S3 objects encrypted with a customer managed AWS KMS key, permissions must be granted in all relevant policy layers. "Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to the account ID of the Strategy account" is required so the Creative account's bucket resource policy trusts principals from the Strategy account for read access. "Update the custom KMS key policy in the Creative account to grant decrypt permissions to the `strategy_reviewer` IAM role" is also required because S3 must be able to use the KMS key to decrypt the objects on behalf of the assumed role. Finally, "Update the `strategy_reviewer` IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key" provides the identity-based permissions that the

assumed role needs. This combination follows AWS's cross-account authorization model: the resource owner must allow access, the external IAM principal must have permissions, and the KMS key policy must permit cross-account cryptographic use. See AWS guidance for [cross-account S3 bucket access](#) and [allowing users in other accounts to use a KMS key](#).

QUESTION NO: 74

A company has migrated an application from on premises to AWS. The application frontend is a static website that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). The application backend is a Python application that runs on three EC2 instances behind another ALB. The EC2 instances are large, general purpose On-Demand Instances that were sized to meet the on-premises specifications for peak usage of the application.

The application averages hundreds of thousands of requests each month. However, the application is used mainly during lunchtime and receives minimal traffic during the rest of the day.

A solutions architect needs to optimize the infrastructure cost of the application without negatively affecting the application availability.

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

- A. Change all the EC2 instances to compute optimized instances that have the same number of cores as the existing EC2 instances.
- B. Move the application frontend to a static website that is hosted on Amazon S3.
- C. Deploy the application frontend by using AWS Elastic Beanstalk. Use the same instance type for the nodes.
- D. Change all the backend EC2 instances to Spot Instances.
- E. Deploy the backend Python application to general purpose burstable EC2 instances that have the same number of cores as the existing EC2 instances.

ANSWER: B E

Explanation:

Move the application frontend to a static website that is hosted on Amazon S3 is correct because the frontend is already static, so running EC2 instances and an Application Load Balancer for that tier adds unnecessary compute and load-balancing cost. Amazon S3 static website hosting provides a highly durable, low-cost hosting model for static assets, and it can be paired with Amazon CloudFront if global caching or HTTPS at the edge is needed. See the AWS guidance for [hosting a static website using Amazon S3](#).

Deploy the backend Python application to general purpose burstable EC2 instances that have the same number of cores as the existing EC2 instances is also correct because the workload has a predictable burst pattern: heavy use around lunchtime and minimal traffic the rest of the day. Burstable general purpose instances are designed for applications that need a baseline level of CPU with the ability to burst when demand increases, often at a lower hourly cost than larger general purpose instances. Keeping the same number of cores helps preserve application capacity while better matching the intermittent utilization pattern. AWS describes this model in the documentation for [burstable performance instances](#).

QUESTION NO: 75

A company is building a serverless application that runs on an AWS Lambda function that is attached to a VPC. The company needs to integrate the application with a new service from an external provider. The external provider supports only requests that come from public IPv4 addresses that are in an allow list.

The company must provide a single public IP address to the external provider before the application can start using the new service.

Which solution will give the application the ability to access the new service?

- A. Deploy a NAT gateway. Associate an Elastic IP address with the NAT gateway. Configure the VPC to use the NAT gateway.

B. Deploy an egress-only internet gateway. Associate an Elastic IP address with the egress-only internet gateway. Configure the elastic network interface on the Lambda function to use the egress-only internet gateway.

C. Deploy an internet gateway. Associate an Elastic IP address with the internet gateway. Configure the Lambda function to use the internet gateway.

D. Deploy an internet gateway. Associate an Elastic IP address with the internet gateway. Configure the default route in the public VPC route table to use the internet gateway.

ANSWER: A

Explanation:

Deploy a NAT gateway. Associate an Elastic IP address with the NAT gateway. Configure the VPC to use the NAT gateway. is correct because a Lambda function connected to a VPC uses elastic network interfaces in the selected subnets and does not automatically receive direct internet access. To make outbound calls to a public third-party service while keeping the Lambda function in private subnets, the private subnet route table should send internet-bound IPv4 traffic, such as 0.0.0.0/0, to a NAT gateway that is deployed in a public subnet.

A NAT gateway uses an Elastic IP address as its public IPv4 address. That Elastic IP provides the stable, single public source address that can be given to the external provider for allow-listing. Outbound connections from the Lambda function are translated by the NAT gateway, so the provider sees requests as coming from that Elastic IP address. AWS documents this as the standard pattern for allowing resources in private subnets to initiate outbound internet connections through a NAT gateway. See [Internet access for VPC-connected Lambda functions](#) and [NAT gateways](#).

QUESTION NO: 76

A company wants to refactor its retail ordering web application that currently has a load-balanced Amazon EC2 instance fleet for web hosting, database API services, and business logic. The company needs to create a decoupled, scalable architecture with a mechanism for retaining failed orders while also minimizing operational costs.

Which solution will meet these requirements?

A. Use Amazon S3 for web hosting with Amazon API Gateway for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use Amazon Elastic Container Service (Amazon ECS) for business logic with Amazon SQS long polling for retaining failed orders.

B. Use AWS Elastic Beanstalk for web hosting with Amazon API Gateway for database API services. Use Amazon MQ for order queuing. Use AWS Step Functions for business logic with Amazon S3 Glacier Deep Archive for retaining failed orders.

C. Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.

D. Use Amazon Lightsail for web hosting with AWS AppSync for database API services. Use Amazon Simple Email Service (Amazon SES) for order queuing. Use Amazon Elastic Kubernetes Service (Amazon EKS) for business logic with Amazon OpenSearch Service for retaining failed orders.

ANSWER: C

Explanation:

Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders. is correct because it replaces the EC2-based application stack with highly managed, serverless services that scale automatically and reduce operational overhead. Amazon S3 can host static web content without requiring servers to manage, patch, or scale. AWS AppSync provides a managed API layer that can connect clients to backend data sources and business workflows, which fits the need to separate database API services from the web tier. Amazon SQS is a durable managed queue that decouples order submission from downstream processing, allowing the system to absorb traffic spikes and process orders asynchronously. AWS Lambda is well suited for event-driven business logic because it can process messages from SQS without provisioning compute capacity. An Amazon SQS dead-letter queue provides the required

mechanism to retain messages that cannot be processed successfully after the configured retry attempts, so failed orders can be investigated and reprocessed later. See [Amazon SQS dead-letter queues](#) and [Using AWS Lambda with Amazon SQS](#).

QUESTION NO: 77

A software company hosts an application on AWS with resources in multiple AWS accounts and Regions. The application runs on a group of Amazon EC2 instances in an application VPC located in the us-east-1 Region with an IPv4 CIDR block of 10.10.0.0/16. In a different AWS account, a shared services VPC is located in the us-east-2 Region with an IPv4 CIDR block of 10.10.10.0/24. When a cloud engineer uses AWS CloudFormation to attempt to peer the application

VPC with the shared services VPC, an error message indicates a peering failure.

Which factors could cause this error? (Choose two.)

- A. The IPv4 CIDR ranges of the two VPCs overlap
- B. The VPCs are not in the same Region
- C. One or both accounts do not have access to an Internet gateway
- D. One of the VPCs was not shared through AWS Resource Access Manager
- E. The IAM role in the peer acceptor account does not have the correct permissions

ANSWER: A E

Explanation:

The IPv4 CIDR ranges of the two VPCs overlap is correct because Amazon VPC peering cannot be created between VPCs that have matching or overlapping IPv4 or IPv6 CIDR blocks. In this scenario, the shared services VPC CIDR block of 10.10.10.0/24 is contained within the application VPC CIDR block of 10.10.0.0/16, so AWS will reject the peering request. This restriction applies regardless of whether the VPCs are in the same account, different accounts, the same Region, or different Regions. See the AWS VPC peering limitations documentation: [Invalid VPC peering connection configurations](#).

The IAM role in the peer acceptor account does not have the correct permissions is also correct for a cross-account VPC peering connection created with AWS CloudFormation. When the requester and acceptor VPCs are in different AWS accounts, CloudFormation requires an IAM role in the acceptor account that the requester can assume to accept the peering connection. If that role is missing or lacks the needed EC2 permissions, the peering operation can fail. AWS documents this requirement for the [AWS::EC2::VPCPeeringConnection](#) resource.

QUESTION NO: 78

A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with the number of files rapidly declining after business hours.

What is the MOST cost-effective migration recommendation?

- A. Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in an Amazon S3 bucket.
- B. Create a queue using Amazon M. Configure the existing web server to publish to the new queue. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the files. Store the processed files in Amazon EFS. Shut down the EC2 instance after the task is complete.
- C. Create a queue using Amazon MO. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in Amazon EFS.

D. Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the files. Scale the EC2 instances based on the SQS queue length. Store the processed files in an Amazon S3 bucket.

ANSWER: D

Explanation:

The recommendation to use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue, scale based on queue length, and store processed files in Amazon S3 is the most cost-effective fit for this workload. The key requirement is asynchronous processing of files that can take up to 1 hour each, with demand that rises during business hours and falls afterward. A queue-based architecture decouples uploads from processing, allowing the application to absorb bursts without requiring processing capacity to be permanently provisioned for peak load. EC2 Auto Scaling can then add or remove worker instances according to backlog metrics, such as queue depth or backlog per instance, so the company pays for compute capacity only when there is work to process. AWS documents this pattern for scaling worker fleets based on Amazon SQS queue demand: [Scaling based on Amazon SQS](#). Amazon S3 is also a strong storage target for processed media files because it provides highly durable, elastic object storage without the need to manage NAS infrastructure: [Amazon S3 User Guide](#).

QUESTION NO: 79

A company has an application that runs on Amazon EC2 instances. A solutions architect is designing VPC infrastructure in an AWS Region where the application needs to access an Amazon Aurora DB cluster. The EC2 instances are all associated with the same security group. The DB cluster is associated with its own security group.

The solutions architect needs to add rules to the security groups to provide the application with least privilege access to the DB cluster.

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

- A.** Add an inbound rule to the EC2 instances ' security group. Specify the DB cluster ' s security group as the source over the default Aurora port.
- B.** Add an outbound rule to the EC2 instances ' security group. Specify the DB cluster ' s security group as the destination over the default Aurora port.
- C.** Add an inbound rule to the DB cluster ' s security group. Specify the EC2 instances ' security group as the source over the default Aurora port.
- D.** Add an outbound rule to the DB cluster ' s security group. Specify the EC2 instances ' security group as the destination over the default Aurora port.
- E.** Add an outbound rule to the DB cluster ' s security group. Specify the EC2 instances ' security group as the destination over the ephemeral ports.

ANSWER: B C

Explanation:

The correct combination is to add an outbound rule to the EC2 instances ' security group. Specify the DB cluster ' s security group as the destination over the default Aurora port and to add an inbound rule to the DB cluster ' s security group. Specify the EC2 instances ' security group as the source over the default Aurora port. This provides least privilege because the application instances are permitted to initiate database connections only to resources that are members of the Aurora DB cluster security group, and the DB cluster accepts database traffic only from resources that are members of the EC2 instances' security group. Security group references are a recommended way to control traffic between resources without relying on changing private IP addresses. Security groups are stateful, so return traffic for an allowed connection is automatically permitted; the key rules for this design are the client-side outbound rule and the database-side inbound rule on the Aurora database port. For Aurora MySQL-compatible clusters the default port is typically 3306, and for Aurora PostgreSQL-compatible clusters the default port is typically 5432, unless the cluster was configured differently. See AWS documentation on [security group rules](#) and [controlling access with security groups for Amazon Aurora](#).

QUESTION NO: 80

A company is migrating a legacy application from an on-premises data center to AWS. The application consists of a single application server and a Microsoft SQL Server database server. Each server is deployed on a VMware VM that consumes 500 TB of data across multiple attached volumes. The company has established a 10 Gbps AWS Direct Connect connection from the closest AWS Region to its on-premises data center. The Direct Connect connection is not currently in use by other services. Which combination of steps should a solutions architect take to migrate the application with the LEAST amount of downtime? (Choose two.)

- A. Use an AWS Server Migration Service (AWS SMS) replication job to migrate the database server VM to AWS.
- B. Use VM Import/Export to import the application server VM.
- C. Export the VM images to an AWS Snowball Edge Storage Optimized device.
- D. Use an AWS Server Migration Service (AWS SMS) replication job to migrate the application server VM to AWS.
- E. Use an AWS Database Migration Service (AWS DMS) replication instance to migrate the database to an Amazon RDS DB instance.

ANSWER: A D

Explanation:

The correct approach is to use AWS Server Migration Service replication jobs for both the database server VM and the application server VM. AWS SMS is designed for lift-and-shift migration of VMware virtual machines to AWS by automating incremental replication of server volumes and creating Amazon Machine Images that can be launched as EC2 instances. Because replication can occur while the source VMs continue to run, the long initial transfer of hundreds of terabytes can happen without requiring the application to be offline. The established 10 Gbps AWS Direct Connect connection provides a high-bandwidth private path for the initial and ongoing replication traffic, helping the migration complete efficiently while preserving the existing server configuration, operating system, application stack, and SQL Server deployment. Downtime can then be limited to the final synchronization and cutover window, when the on-premises application is stopped, the last changes are replicated, and the AWS-based instances are launched. This matches the requirement to move the existing two-VM application with the least downtime. See [AWS Server Migration Service](#) and the [AWS Direct Connect User Guide](#) for related AWS service details.

QUESTION NO: 81

A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.

A solutions architect has created an IAM role that is named `strategy_reviewer` in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Account.

The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to the account ID of the Strategy account
- B. Update the `strategy_reviewer` IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.
- C. Update the custom KMS key policy in the Creative account to grant decrypt permissions to the `strategy_reviewer` IAM role.

D. Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to an anonymous user.

E. Update the custom KMS key policy in the Creative account to grant encrypt permissions to the strategy_reviewer IAM role.

F. Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key

ANSWER: A C F

Explanation:

For cross-account access to Amazon S3 objects that are encrypted with a customer managed AWS KMS key, permissions must be granted in both accounts and for both services involved. Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to the account ID of the Strategy account is correct because the bucket owner must allow the external Strategy account to read the S3 objects. Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key is also required because the assumed role needs an identity-based policy that allows the specific S3 read actions and the KMS decrypt action. Finally, Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role is correct because AWS KMS key policies control cross-account use of customer managed keys; the external role cannot decrypt the S3 objects unless the key policy permits it. Together, these steps provide the minimum access path needed for viewing KMS-encrypted objects across accounts. See AWS guidance on [cross-account S3 AccessDenied troubleshooting](#) and the AWS KMS documentation on [allowing users in other accounts to use a KMS key](#).

QUESTION NO: 82

A company is changing the way that it handles patching of Amazon EC2 instances in its application account. The company currently patches instances over the internet by using a NAT gateway in a VPC in the application account. The company has EC2 instances set up as a patch source repository in a dedicated private VPC in a core account. The company wants to use AWS Systems Manager Patch Manager and the patch source repository in the core account to patch the EC2 instances in the application account. The company must prevent all EC2 instances in the application account from accessing the internet. The EC2 instances in the application account need to access Amazon S3, where the application data is stored. These EC2 instances need connectivity to Systems Manager and to the patch source repository in the private VPC in the core account. Which solution will meet these requirements?

A. Create a network ACL that blocks outbound traffic on port 80. Associate the network ACL with all subnets in the application account. In the application account and the core account, deploy one EC2 instance that runs a custom VPN server. Create a VPN tunnel to access the private VPC. Update the route table in the application account.

B. Create private VIFs for Systems Manager and Amazon S3. Delete the NAT gateway from the VPC in the application account. Create a transit gateway to access the patch source repository EC2 instances in the core account. Update the route table in the core account.

C. Create VPC endpoints for Systems Manager and Amazon S3. Delete the NAT gateway from the VPC in the application account. Create a VPC peering connection to access the patch source repository EC2 instances in the core account. Update the route tables in both accounts.

D. Create a network ACL that blocks inbound traffic on port 80. Associate the network ACL with all subnets in the application account. Create a transit gateway to access the patch source repository EC2 instances in the core account. Update the route tables in both accounts.

ANSWER: C

Explanation:

Create VPC endpoints for Systems Manager and Amazon S3. Delete the NAT gateway from the VPC in the application account. Create a VPC peering connection to access the patch source repository EC2 instances in the core account. Update the route tables in both accounts is correct because it removes the internet egress path while preserving all required private connectivity. AWS Systems Manager managed instances can communicate privately with Systems Manager by using AWS PrivateLink interface VPC endpoints, typically including endpoints for Systems Manager, EC2 messages, and Systems

Manager messages. Amazon S3 access can be provided without internet access by using an S3 gateway endpoint or interface endpoint, depending on the design. This allows the application instances to reach S3 while avoiding a NAT gateway or internet gateway path. For the private patch repository hosted on EC2 instances in the core account VPC, VPC peering provides private IP connectivity between the application VPC and core VPC, as long as route tables and security controls allow the traffic. Together, these components satisfy the requirement to block internet access while enabling Patch Manager operations, S3 access, and connectivity to the internal patch source repository. See AWS guidance for [using VPC endpoints with Systems Manager](#) and [VPC peering](#).

QUESTION NO: 83

A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53.

A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.

Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

- A. Create a dynamic webpage that runs on an Amazon EC2 instance. Configure the webpage to use the JSON document in combination with the event message to look up and respond with a redirect URL.
- B. Create an Application Load Balancer that includes HTTP and HTTPS listeners.
- C. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.
- D. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.
- E. Create an Amazon CloudFront distribution. Deploy a Lambda@Edge function.
- F. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.

ANSWER: C E F

Explanation:

The correct combination is to create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL, create an Amazon CloudFront distribution and deploy a Lambda@Edge function, and create an SSL certificate by using AWS Certificate Manager with the domains as Subject Alternative Names. Lambda@Edge is well suited for this use case because it can run logic at CloudFront edge locations in response to viewer requests, inspect request attributes such as the Host header, and immediately return an HTTP redirect response without operating servers. CloudFront provides a managed global entry point that can accept both HTTP and HTTPS requests for the marketing domains. For HTTPS, CloudFront requires a trusted certificate that covers the alternate domain names configured on the distribution, and ACM certificates for CloudFront must be requested in the US East (N. Virginia) Region. Including all domains as Subject Alternative Names allows one certificate to support the 10 domain names. This approach minimizes operational effort because it avoids managing compute instances or load balancers while using Route 53 alias records to point the domains to CloudFront. See [Lambda@Edge documentation](#) and [CloudFront alternate domain name HTTPS requirements](#).

QUESTION NO: 84

A company is building a hybrid environment that includes servers in an on-premises data center and in the AWS Cloud. The company has deployed Amazon EC2 instances in three VPCs. Each VPC is in a different AWS Region. The company has established an AWS Direct Connect connection to the data center from the Region that is closest to the data center.

The company needs the servers in the on-premises data center to have access to the EC2 instances in all three VPCs. The servers in the on-premises data center also must have access to AWS public services.

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

- A. Create a Direct Connect gateway in the Region that is closest to the data center. Attach the Direct Connect connection to the Direct Connect gateway. Use the
- B. Direct Connect gateway to connect the VPCs in the other two Regions.
- C. Set up additional Direct Connect connections from the on-premises data center to the other two Regions.
- D. Create a private VIF. Establish an AWS Site-to-Site VPN connection over the private VIF to the VPCs in the other two Regions.
- E. Create a public VIF. Establish an AWS Site-to-Site VPN connection over the public VIF to the VPCs in the other two Regions.
- F. Use VPC peering to establish a connection between the VPCs across the Regions. Create a private VIF with the existing Direct Connect connection to connect to the peered VPCs.

ANSWER: A E

Explanation:

Create a Direct Connect gateway in the Region that is closest to the data center is correct because a Direct Connect gateway is designed to let an existing AWS Direct Connect connection reach VPCs across multiple AWS Regions through private connectivity. This avoids provisioning separate Direct Connect connections to each Region and lets the company reuse the existing connection for the hybrid architecture. AWS documents that Direct Connect gateways can be associated with virtual private gateways or transit gateways in multiple Regions, which fits the requirement to reach EC2 instances in VPCs that are spread across Regions. See the AWS Direct Connect gateway documentation at [AWS Direct Connect gateways](#).

Create a public VIF. Establish an AWS Site-to-Site VPN connection over the public VIF to the VPCs in the other two Regions is also correct because a public virtual interface provides access to AWS public services over Direct Connect. Public VIFs advertise AWS public IP prefixes, which enables private data center connectivity to public AWS service endpoints without using the internet path. AWS describes public virtual interfaces in the Direct Connect virtual interface documentation at [AWS Direct Connect virtual interfaces](#).

QUESTION NO: 85

A company runs a microservice as an AWS Lambda function. The microservice writes data to an on-premises SQL database that supports a limited number of concurrent connections. When the number of Lambda function invocations is too high, the database crashes and causes application downtime. The company has an AWS Direct Connect connection between the company's VPC and the on-premises data center. The company wants to protect the database from crashes. Which solution will meet these requirements?

- A. Write the data to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the Lambda function to read from the queue and write to the existing database. Set a reserved concurrency limit on the Lambda function that is less than the number of connections that the database supports.
- B. Create a new Amazon Aurora Serverless DB cluster. Use AWS DataSync to migrate the data from the existing database to Aurora Serverless. Reconfigure the Lambda function to write to Aurora.
- C. Create an Amazon RDS Proxy DB instance. Attach the RDS Proxy DB instance to the Amazon RDS DB instance. Reconfigure the Lambda function to write to the RDS Proxy DB instance.
- D. Write the data to an Amazon Simple Notification Service (Amazon SNS) topic. Invoke the Lambda function to write to the existing database when the topic receives new messages. Configure provisioned concurrency for the Lambda function to be equal to the number of connections that the database supports.

ANSWER: A

Explanation:

Writing the data to an Amazon Simple Queue Service (Amazon SQS) queue and configuring the Lambda function to consume from that queue is the correct approach because it decouples request ingestion from database writes. SQS can absorb bursts of incoming messages, while Lambda processes those messages at a controlled rate. Setting reserved concurrency on the Lambda function to a value lower than the number of database connections supported by the on-premises SQL database directly limits the maximum number of concurrent Lambda executions that can attempt database connections. This protects the database from connection storms while still allowing the application to accept incoming work reliably through the queue.

Reserved concurrency is specifically designed to set both the maximum and guaranteed concurrency for a Lambda function, making it appropriate for throttling access to constrained downstream systems. Lambda event source mappings for SQS also support queued, retryable processing, which helps preserve data during temporary database slowdowns. For more details, see the AWS documentation for [configuring reserved concurrency for Lambda](#) and [using Lambda with Amazon SQS](#).

QUESTION NO: 86

A retail company is mounting IoT sensors in all of its stores worldwide. During the manufacturing of each sensor, the company's private certificate authority (CA) issues an X.509 certificate that contains a unique serial number. The company then deploys each certificate to its respective sensor. A solutions architect needs to give the sensors the ability to send data to AWS after they are installed. Sensors must not be able to send data to AWS until they are installed. Which solution will meet these requirements?

- A.** Create an AWS Lambda function that can validate the serial number. Create an AWS IoT Core provisioning template. Include the SerialNumber parameter in the Parameters section. Add the Lambda function as a pre-provisioning hook. During manufacturing, call the RegisterThing API operation and specify the template and parameters.
- B.** Create an AWS Step Functions state machine that can validate the serial number. Create an AWS IoT Core provisioning template. Include the SerialNumber parameter in the Parameters section. Specify the Step Functions state machine to validate parameters. Call the StartThingRegistrationTask API operation during installation.
- C.** Create an AWS Lambda function that can validate the serial number. Create an AWS IoT Core provisioning template. Include the SerialNumber parameter in the Parameters section. Add the Lambda function as a pre-provisioning hook. Register the CA with AWS IoT Core, specify the provisioning template, and set the allow-auto-registration parameter.
- D.** Create an AWS IoT Core provisioning template. Include the SerialNumber parameter in the Parameters section. Include parameter validation in the template. Provision a claim certificate and a private key for each device that uses the CA. Grant AWS IoT Core service permissions to update AWS IoT things during provisioning.

ANSWER: C

Explanation:

Registering the company's private CA with AWS IoT Core, associating an AWS IoT provisioning template, enabling auto-registration, and using a Lambda pre-provisioning hook is the correct approach. AWS IoT Core supports just-in-time provisioning for devices that already have certificates signed by a registered CA. When a sensor first connects to AWS IoT Core after installation, AWS IoT can automatically register the device certificate, create or update the AWS IoT thing, and attach the required policy based on the provisioning template. Until that first connection and successful provisioning workflow occurs, the certificate is not an active, authorized AWS IoT identity that can publish data.

The Lambda pre-provisioning hook is important because it lets the company validate device-specific attributes, such as the unique serial number from the X.509 certificate or provisioning parameters, before AWS IoT completes provisioning. This allows only legitimate installed sensors with approved serial numbers to become active. AWS documents this pattern as fleet provisioning and just-in-time provisioning with pre-provisioning hooks for additional validation logic. See [AWS IoT Core just-in-time provisioning](#) and [AWS IoT pre-provisioning hooks](#).

QUESTION NO: 87

A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances. To meet regulatory and business requirements, the company must make the following changes for data backups:

- Backups must be retained based on custom daily, weekly, and monthly requirements.
- Backups must be replicated to at least one other AWS Region immediately after capture.
- The backup solution must provide a single source of backup status across the AWS

environment. • The backup solution must send immediate notifications upon failure of any resource backup. Which combination of steps will meet these requirements with the LEAST amount of operational overhead? (Choose three.)

- A. Create an AWS Backup plan with a backup rule for each of the retention requirements.
- B. Configure an AWS Backup plan to copy backups to another Region.
- C. Create an AWS Lambda function to replicate backups to another Region and send notification if a failure occurs.
- D. Add an Amazon Simple Notification Service (Amazon SNS) topic to the backup plan to send a notification for finished jobs that have any status except `BACKUP_JOB_COMPLETED`.
- E. Create an Amazon Data Lifecycle Manager (Amazon DLM) snapshot lifecycle policy for each of the retention requirements.
- F. Set up RDS snapshots on each database.

ANSWER: A B D

Explanation:

AWS Backup is the best fit because it provides a centralized, managed backup service for supported AWS resources, including Amazon EC2, Amazon EFS, and Amazon RDS. Creating an AWS Backup plan with a backup rule for each of the retention requirements allows the company to define separate daily, weekly, and monthly schedules and lifecycle retention settings without building custom automation. Configuring an AWS Backup plan to copy backups to another Region satisfies the requirement to replicate backups immediately after capture by using managed cross-Region copy actions. This also keeps the design operationally simple because the replication behavior is part of the backup policy rather than a separate workflow. Adding an Amazon Simple Notification Service topic to send notifications for backup jobs that do not complete successfully provides prompt visibility into failed or otherwise unsuccessful backup activity. Together, these steps use AWS Backup as the single source for backup configuration, monitoring, status, and compliance across the environment. For more details, see the AWS Backup documentation on [centralized backup management](#) and [cross-Region backup copies](#).

QUESTION NO: 88

A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode.

A recent RDS database failover test caused a 40-second outage to the application. A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Use Amazon Elasti Cache for Memcached in front of the database
- B. Use Amazon Elasti Cache for Redis in front of the database.
- C. Use RDS Proxy in front of the database
- D. Migrate the database to Amazon Aurora MySQL
- E. Create an Amazon Aurora Replica
- F. Create an RDS for MySQL read replica

ANSWER: C D E

Explanation:

The correct combination is to use RDS Proxy in front of the database, migrate the database to Amazon Aurora MySQL, and create an Amazon Aurora Replica. Aurora MySQL separates compute from storage and uses a distributed, fault-tolerant storage layer across multiple Availability Zones, which helps Aurora perform failover more quickly than a traditional Amazon RDS for MySQL Multi-AZ DB instance. An Amazon Aurora Replica in the same Aurora DB cluster can be promoted automatically during failover, giving the cluster a ready compute target to assume the writer role. AWS states that Aurora is

designed for high availability and typically completes failover in less than 30 seconds, and proper replica configuration is a key part of that design. Placing RDS Proxy in front of the database further reduces application impact during failover by pooling and preserving connections, quickly routing traffic to the new writer, and minimizing the time the application spends reconnecting. AWS specifically notes that RDS Proxy can make failovers faster and more resilient for applications. Together, these steps address both database failover speed and application connection recovery, which is necessary to reduce the observed outage below 20 seconds. See [Amazon Aurora high availability](#) and [Amazon RDS Proxy](#).

QUESTION NO: 89

A retail company needs to provide a series of data files to another company, which is its business partner. These files are saved in an Amazon S3 bucket under Account A.

which belongs to the retail company. The business partner company wants one of its IAM users, User_DataProcessor, to access the files from its own AWS account (Account B).

Which combination of steps must the companies take so that User_DataProcessor can access the S3 bucket successfully? (Select TWO.)

- A. Turn on the cross-origin resource sharing (CORS) feature for the S3 bucket in Account
- B. In Account A, set the S3 bucket policy to the following:
- C. In Account A, set the S3 bucket policy to the following:
- D. In Account B, set the permissions of User_DataProcessor to the following:
- E. In Account B, set the permissions of User_DataProcessor to the following:

ANSWER: C D

Explanation:

The correct combination is to configure access on both sides of the cross-account relationship: the S3 bucket owner in Account A must add a bucket policy that allows the appropriate Account B principal to access the required bucket resources, and Account B must grant User_DataProcessor IAM permissions to perform the required S3 actions. For cross-account access to Amazon S3, an allow in only one account is not enough in the standard model: the resource-owning account must trust the external principal through the bucket policy, and the principal's own account must authorize the user through an identity-based IAM policy. This lets User_DataProcessor call actions such as listing the bucket or reading objects, depending on the permissions granted in both policies. The effective access is the intersection of what the bucket policy in Account A permits and what the IAM policy for User_DataProcessor in Account B permits. AWS documents this pattern for granting cross-account S3 bucket access and recommends using resource-based bucket policies together with IAM permissions in the trusted account. See [AWS re:Post Knowledge Center: cross-account access to Amazon S3](#) and [Amazon S3 documentation: bucket owner granting cross-account bucket permissions](#).

QUESTION NO: 90

A company wants to send data from its on-premises systems to Amazon S3 buckets. The company created the S3 buckets in three different accounts. The company must send the data privately without the data traveling across the internet. The company has no existing dedicated connectivity to AWS.

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

- A. Establish a networking account in the AWS Cloud. Create a private VPC in the networking account. Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC.
- B. Establish a networking account in the AWS Cloud. Create a private VPC in the networking account. Set up an AWS Direct Connect connection with a public VIF between the on-premises environment and the private VPC.
- C. Create an Amazon S3 interface endpoint in the networking account.
- D. Create an Amazon S3 gateway endpoint in the networking account.

E. Establish a networking account in the AWS Cloud Create a private VPC in the networking account. Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account.

ANSWER: A C

Explanation:

Establish a networking account in the AWS Cloud Create a private VPC in the networking account. Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC. is correct together with Create an Amazon S3 interface endpoint in the networking account. AWS Direct Connect provides dedicated private connectivity from the on-premises environment to AWS, and a private virtual interface connects that dedicated circuit to a VPC through a virtual private gateway or Direct Connect gateway. To reach Amazon S3 using private IP addressing from on premises, the VPC can host an interface VPC endpoint for Amazon S3, which is powered by AWS PrivateLink. That endpoint creates elastic network interfaces with private IP addresses in the VPC, so on-premises systems can send S3 traffic over the Direct Connect private VIF to those private endpoint addresses rather than across the public internet. The S3 buckets can be in different AWS accounts because access is controlled through S3 bucket policies, IAM permissions, and the interface endpoint policy. This design centralizes the private network entry point in the networking account while still allowing controlled access to S3 buckets in the other accounts. See [AWS PrivateLink for Amazon S3](#) and [AWS Direct Connect virtual interfaces](#).

QUESTION NO: 91

A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS for PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.

What changes to the current architecture will reduce operational overhead and support the product release?

- A.** Create an EC2 Auto Scaling group behind an Application Load Balancer. Create additional read replicas for the DB instance. Create Amazon Kinesis data streams and configure the application services to use the data streams. Store and serve static content directly from Amazon S3.
- B.** Create an EC2 Auto Scaling group behind an Application Load Balancer. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. Create Amazon Kinesis data streams and configure the application services to use the data streams. Store and serve static content directly from Amazon S3.
- C.** Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balancer. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
- D.** Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

ANSWER: D

Explanation:

Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution. is correct because it replaces self-managed compute and messaging components with managed AWS services while improving scalability for the expected traffic spike. Amazon EKS with AWS Fargate lets the company run containers without managing worker nodes, patching EC2 instances, or handling cluster capacity at the instance level. With autoscaling and an Application Load Balancer, the application tier can scale horizontally and receive traffic through a managed, highly available entry point. Amazon RDS read replicas help scale read-heavy database workloads, which is commonly needed during order surges and product launches. Amazon Managed Streaming for

Apache Kafka reduces the operational burden of running Kafka brokers on EC2 by handling broker provisioning, patching, and high availability. Finally, serving static content from Amazon S3 through Amazon CloudFront reduces load on the application, improves global performance, and absorbs spikes at the edge. These choices align with AWS managed-service best practices for lowering operations while increasing elasticity. References: [AWS Fargate with Amazon EKS](#) and [Amazon Managed Streaming for Apache Kafka](#).

QUESTION NO: 92

A company is running an application in the AWS Cloud. The company's security team must approve the creation of all new IAM users. When a new IAM user is created, all access for the user must be removed automatically. The security team must then receive a notification to approve the user. The company has a multi-Region AWS CloudTrail trail in the AWS account. Which combination of steps will meet these requirements? (Choose three.)

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule. Define a pattern with the detail-type value set to AWS API Call via CloudTrail and an eventName of CreateUser.
- B. Configure CloudTrail to send a notification for the CreateUser event to an Amazon Simple Notification Service (Amazon SNS) topic.
- C. Invoke a container that runs in Amazon Elastic Container Service (Amazon ECS) with AWS Fargate technology to remove access.
- D. Invoke an AWS Step Functions state machine to remove access.
- E. Use Amazon Simple Notification Service (Amazon SNS) to notify the security team.
- F. Use Amazon Pinpoint to notify the security team.

ANSWER: A D E

Explanation:

The correct combination is to create an Amazon EventBridge rule that matches CloudTrail management events for CreateUser, invoke an AWS Step Functions state machine to remove access, and use Amazon Simple Notification Service to notify the security team. IAM CreateUser is recorded by CloudTrail as a management event, and EventBridge can match AWS API calls delivered through CloudTrail by using an event pattern such as detail-type "AWS API Call via CloudTrail" and eventName "CreateUser." This provides near-real-time automation whenever a new IAM user is created. See the AWS documentation for using [EventBridge with CloudTrail events](#). A Step Functions state machine is appropriate for the remediation workflow because it can orchestrate multiple AWS API actions, including IAM-related SDK integrations, to remove access such as credentials, permissions, or login profiles before approval. AWS documents these service integrations in [Step Functions AWS SDK integrations](#). After the automated remediation runs, Amazon SNS is a standard AWS service for fan-out notifications to email, HTTPS endpoints, or other subscribers, making it suitable for alerting the security team for approval.

QUESTION NO: 93

A company has an Amazon VPC that is divided into a public subnet and a private subnet. A web application runs in Amazon VPC. and each subnet has its own NACL. The public subnet has a CIDR of 10.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet.

Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets.

What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

- A. An inbound rule for port 80 from source 0.0.0.0/0
- B. An inbound rule for port 80 from source 10.0.0/24
- C. An outbound rule for port 80 to destination 0.0.0.0/0

D. An outbound rule for port 80 to destination 10.0.0.0/24

E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

ANSWER: B E

Explanation:

The private subnet network ACL must allow the traffic that reaches the EC2 web servers and the return traffic back to the Application Load Balancer. Network ACLs are stateless, so allowing the request in one direction does not automatically allow the response in the opposite direction. Therefore, **An inbound rule for port 80 from source 10.0.0.0/24** is correct because the ALB resides in the public subnet, and it forwards HTTP requests to targets on port 80 in the private subnet. The rule is scoped to the public subnet CIDR, which satisfies the requirement to allow only the necessary subnet-to-subnet traffic. **An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24**

is also correct because the EC2 instances' response traffic returns to the ALB using ephemeral ports on the client side. AWS documents that NACLs require explicit inbound and outbound rules and that ephemeral port ranges must be considered for return traffic. See the AWS guidance for [network ACLs](#) and [custom network ACL rules and ephemeral ports](#).

QUESTION NO: 94

A company has multiple AWS accounts that are in an organization in AWS Organizations. The company needs to store AWS account activity and query the data from a central location by using SQL.

Which solution will meet these requirements?

A. Create an AWS CloudTrail trail in each account. Specify CloudTrail management events for the trail. Configure CloudTrail to send the events to Amazon CloudWatch Logs. Configure CloudWatch cross-account observability. Query the data in CloudWatch Logs Insights.

B. Use a delegated administrator account to create an AWS CloudTrail Lake data store. Specify CloudTrail management events for the data store. Enable the data store for all accounts in the organization. Query the data in CloudTrail Lake.

C. Use a delegated administrator account to create an AWS CloudTrail trail. Specify CloudTrail management events for the trail. Enable the trail for all accounts in the organization. Keep all other settings as default. Query the CloudTrail data from the CloudTrail event history page.

D. Use AWS CloudFormation StackSets to deploy AWS CloudTrail Lake data stores in each account. Specify CloudTrail management events for the data stores. Keep all other settings as default. Query the data in CloudTrail Lake.

ANSWER: B

Explanation:

Use a delegated administrator account to create an AWS CloudTrail Lake data store. Specify CloudTrail management events for the data store. Enable the data store for all accounts in the organization. Query the data in CloudTrail Lake. is correct because CloudTrail Lake is designed for centralized, long-term storage and SQL-based analysis of CloudTrail events. In an AWS Organizations environment, an organization event data store can collect events from all accounts in the organization, giving the company a single place to retain and analyze account activity. CloudTrail management events capture control plane activity such as API calls, console sign-ins, and changes to AWS resources, which aligns with the requirement to store AWS account activity. CloudTrail Lake also provides a managed query capability that uses SQL, so the company can query the collected event data directly without building and operating a separate analytics pipeline. A delegated administrator account can be used to manage CloudTrail Lake for the organization, which supports centralized governance while avoiding the need to configure independent query stores in every account. For more detail, see the AWS documentation for [AWS CloudTrail Lake](#) and [CloudTrail management events](#).

QUESTION NO: 95

A company is collecting a large amount of data from a fleet of IoT devices. Data is stored as Optimized Row Columnar (ORC) files in the Hadoop Distributed File System (HDFS) on a persistent Amazon EMR cluster. The company's data

analytics team queries the data by using SQL in Apache Presto deployed on the same EMR cluster. Queries scan large amounts of data, always run for less than 15 minutes, and run only between 5 PM and 10 PM.

The company is concerned about the high cost associated with the current solution. A solutions architect must propose the most cost-effective solution that will allow SQL data queries.

Which solution will meet these requirements?

- A. Store data in Amazon S3. Use Amazon Redshift Spectrum to query data.
- B. Store data in Amazon S3. Use the AWS Glue Data Catalog and Amazon Athena to query data.
- C. Store data in EMR File System (EMRFS). Use Presto on Amazon EMR to query data.
- D. Store data in Amazon Redshift. Use Amazon Redshift to query data.

ANSWER: B

Explanation:

Store data in Amazon S3. Use the AWS Glue Data Catalog and Amazon Athena to query data is the most cost-effective solution for this workload. Amazon Athena is a serverless interactive query service that lets users run standard SQL directly against data stored in Amazon S3, so the company does not need to keep a persistent EMR cluster running when queries are only needed during a limited daily time window. Athena supports columnar formats such as ORC, which is important because the existing data is already stored in ORC files and queries scan large volumes of data. Using ORC with Athena can reduce the amount of data scanned, which directly reduces query cost because Athena pricing is based primarily on the amount of data scanned per query. The AWS Glue Data Catalog provides the metadata layer that Athena uses to define tables, schemas, and partitions over the S3 data. This combination is well suited for intermittent SQL analytics workloads where maintaining always-on compute capacity would be unnecessarily expensive. See [Amazon Athena documentation](#) and [Athena columnar storage formats](#).

QUESTION NO: 96

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? (Choose 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.

ANSWER: A D F

Explanation:

Configure AWS CloudTrail to log S3 data events is correct because S3 data events provide object-level auditing for API activity such as object reads, writes, and deletes. This is the appropriate AWS-native audit mechanism when a company must log activity against objects in an S3 bucket. 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 is correct because S3 can integrate with EventBridge for object-level events, and EventBridge can route matching deletion events to SNS, where the security team can subscribe by email. Configure a new S3 bucket to store the logs with an S3 Lifecycle policy is correct because CloudTrail logs are commonly delivered to S3, and lifecycle rules can transition older logs to lower-cost storage classes and retain them for the required 5-year period cost-effectively. Together, these steps provide object activity logging,

deletion notifications, and long-term retention with storage cost optimization. Relevant AWS references include [Logging Amazon S3 API calls using AWS CloudTrail](#) and [Using EventBridge with Amazon S3 events](#).

QUESTION NO: 97

An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list. The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets. How should a solutions architect ensure that the web application can continue to call the third-party API after the migration?

- A. Associate a block of customer-owned public IP addresses to the VPC. Enable public IP addressing for public subnets in the VPC.
- B. Register a block of customer-owned public IP addresses in the AWS account. Create Elastic IP addresses from the address block and assign them to the NAT gateways in the VPC.
- C. Create Elastic IP addresses from the block of customer-owned IP addresses. Assign the static Elastic IP addresses to the ALB.
- D. Register a block of customer-owned public IP addresses in the AWS account. Set up AWS Global Accelerator to use Elastic IP addresses from the address block. Set the ALB as the accelerator endpoint.

ANSWER: B

Explanation:

Registering a block of customer-owned public IP addresses in the AWS account, then creating Elastic IP addresses from that address block and assigning them to the NAT gateways in the VPC, is the correct approach. The web application instances are in private subnets, so their outbound internet traffic to the third-party API will egress through the NAT gateways. A public NAT gateway uses its associated Elastic IP address as the source public IP for outbound connections. By bringing the existing customer-owned public CIDR block to AWS using Bring Your Own IP (BYOIP), the customer can continue using the same public address range that the third party already allows through its firewall. Creating Elastic IP addresses from that BYOIP pool and attaching them to the NAT gateways ensures the third-party API sees traffic sourced from the approved CIDR block after migration. This preserves the existing allow-list model while keeping the EC2 instances private and using the intended NAT-based outbound architecture. AWS documents support BYOIP for bringing publicly routable address ranges into AWS and using those addresses as Elastic IPs, and NAT gateways support Elastic IP association for internet-bound traffic. References: [Bring your own IP addresses to Amazon EC2](#) and [NAT gateways](#).

QUESTION NO: 98

A company is using multiple AWS accounts. The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A. The company's applications and databases are running in Account B.

A solutions architect will deploy a two-net application in a new VPC. To simplify the configuration, the db.example.com CNAME record set for the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.

During deployment, the application failed to start. Troubleshooting revealed that db.example.com is not resolvable on the Amazon EC2 instance. The solutions architect confirmed that the record set was created correctly in Route 53.

Which combination of steps should the solutions architect take to resolve this issue? (Select TWO.)

- A. Deploy the database on a separate EC2 instance in the new VPC. Create a record set for the instance's private IP in the private hosted zone.
- B. Use SSH to connect to the application tier EC2 instance. Add an RDS endpoint IP address to the /etc/resolv.conf file.
- C. Create an authorization to associate the private hosted zone in Account A with the new VPC in Account B.
- D. Create a private hosted zone for the example.com domain in Account B. Configure Route 53 replication between AWS accounts.

E. Associate a new VPC in Account B with a hosted zone in Account A. Delete the association authorization In Account A.

ANSWER: C E

Explanation:

The correct steps are to create an authorization to associate the private hosted zone in Account A with the new VPC in Account B, and then associate the new VPC in Account B with the hosted zone in Account A. A Route 53 private hosted zone resolves DNS records only for VPCs that are associated with that hosted zone. Because the hosted zone is in one AWS account and the application EC2 instances are in a VPC in another account, Route 53 requires a cross-account association workflow. First, the hosted zone owner authorizes the VPC association. Then, the VPC owner creates the association between the VPC and the private hosted zone. After the association exists, instances in the VPC can resolve records such as the db.example.com CNAME through the VPC-provided DNS resolver, assuming normal VPC DNS settings are enabled. AWS also recommends deleting the association authorization after the association is complete because the authorization is no longer needed. See the AWS documentation for [associating VPCs and private hosted zones across accounts](#) and the [CreateVPCAssociationAuthorization API](#).

QUESTION NO: 99

A solutions architect is designing an application to accept timesheet entries from employees on their mobile devices. Timesheets will be submitted weekly, with most of the submissions occurring on Friday. The data must be stored in a format that allows payroll administrators to run monthly reports. The infrastructure must be highly available and scale to match the rate of incoming data and reporting requests.

Which combination of steps meets these requirements while minimizing operational overhead? (Select TWO.)

- A.** Deploy the application to Amazon EC2 On-Demand Instances With load balancing across multiple Availability Zones. Use scheduled Amazon EC2 Auto Scaling to add capacity before the high volume of submissions on Fridays.
- B.** Deploy the application in a container using Amazon Elastic Container Service (Amazon ECS) with load balancing across multiple Availability Zones. Use scheduled Service Auto Scaling to add capacity before the high volume of submissions on Fridays.
- C.** Deploy the application front end to an Amazon S3 bucket served by Amazon CloudFront. Deploy the application backend using Amazon API Gateway with an AWS Lambda proxy integration.
- D.** Store the timesheet submission data in Amazon Redshift. Use Amazon QuickSight to generate the reports using Amazon Redshift as the data source.
- E.** Store the timesheet submission data in Amazon S3. Use Amazon Athena and Amazon QuickSight to generate the reports using Amazon S3 as the data source.

ANSWER: C E

Explanation:

Deploying the application front end to an Amazon S3 bucket served by Amazon CloudFront and deploying the backend with Amazon API Gateway and an AWS Lambda proxy integration is correct because it provides a highly available, serverless architecture that automatically scales with request volume. This is well suited to a mobile timesheet application with predictable weekly spikes because API Gateway and Lambda can absorb bursts of submissions without requiring instance provisioning or manual capacity management. CloudFront improves global performance and availability for the static front end while S3 provides durable, low-operations hosting for static assets. See [Using AWS Lambda with Amazon API Gateway](#).

Storing the timesheet submission data in Amazon S3 and using Amazon Athena and Amazon QuickSight to generate reports is also correct because S3 is durable, scalable object storage, Athena can query structured data in S3 using standard SQL without managing servers, and QuickSight can visualize/report on those query results. This combination supports monthly reporting while minimizing operational overhead and scaling naturally as data volume grows. See [What is Amazon Athena?](#)

QUESTION NO: 100

A company has an application that stores data in a single Amazon S3 bucket. The company must keep all data for 1 year. The company's security team is concerned that an attacker could gain access to the AWS account through leaked long-term credentials. Which solution will ensure that existing and future objects in the S3 bucket are protected?

- A.** Create a new AWS account that is accessible only to the security team through an assumed role. Create an S3 bucket in the new account. Enable S3 Versioning and S3 Object Lock. Configure a default retention period of 1 year. Set up replication from the existing S3 bucket to the new S3 bucket. Create an S3 Batch Replication job to copy all existing data.
- B.** Use the s3-bucket-versioning-enabled AWS Config managed rule. Configure an automatic remediation action that uses an AWS Lambda function to enable S3 Versioning and MFA Delete on noncompliant resources. Add an S3 Lifecycle rule to delete objects after 1 year.
- C.** Explicitly deny bucket creation from all users and roles except for an AWS Service Catalog launch constraint role. Define a Service Catalog product for the creation of the S3 bucket to force S3 Versioning and MFA Delete to be enabled. Authorize users to launch the product when they need to create an S3 bucket.
- D.** Enable Amazon GuardDuty with the S3 protection feature for the account and the AWS Region. Add an S3 Lifecycle rule to delete objects after 1 year.

ANSWER: A

Explanation:

Creating a new AWS account controlled by the security team, enabling S3 Versioning and S3 Object Lock on a destination bucket, configuring a 1-year default retention period, and replicating both existing and future objects is the correct solution. S3 Object Lock provides write-once-read-many protection by preventing protected object versions from being deleted or overwritten during the retention period. When configured with an appropriate retention period, it directly satisfies the requirement to keep data for 1 year. Placing the protected copy in a separate AWS account that is accessible only through a security-team-controlled assumed role also reduces the risk that leaked long-term credentials from the application account could be used to tamper with the protected data. Replication handles new objects going forward, while S3 Batch Replication is designed to replicate objects that already exist in the source bucket, ensuring both current and future data are covered. This combination provides durable, account-isolated, retention-enforced protection aligned with AWS best practices for immutable S3 storage. See the AWS documentation for [S3 Object Lock](#) and [S3 Batch Replication](#).

QUESTION NO: 101

A company is planning to migrate its applications from an on-premises data center to AWS. The on-premises data center has an AWS Direct Connect connection. The company needs to test IPv6 connectivity in the VPC so that the applications can communicate with more customers worldwide.

A solutions architect has created a VPC with an IPv6 CIDR block.

Which networking configurations will meet these requirements? (Select TWO.)

- A.** Launch an Amazon EC2 instance into a public subnet. Associate an IPv6 address with the instance during launch. Configure a security group, a network ACL, and route tables for IPv6 communication. Associate a virtual private gateway in the VPC with a Direct Connect gateway.
- B.** Launch an Amazon EC2 instance into a private subnet. Associate an IPv6 address with the instance during launch. Configure a security group, a network ACL, and route tables for IPv6 communication. Create a route that directs all IPv6 traffic from the private subnet to a NAT gateway.
- C.** Launch an Amazon EC2 instance into a public subnet. Associate an IPv6 address with the instance during launch. Configure a security group, a network ACL, and route tables for IPv6 communication. Create a route that directs all IPv6 traffic from the public subnet to an internet gateway.
- D.** Launch an Amazon EC2 instance into a private subnet. Associate an IPv6 address with the instance during launch. Configure a security group, a network ACL, and route tables for IPv6 communication. Create a route that directs all IPv6 traffic from the private subnet to a NAT instance.
- E.** Launch an Amazon EC2 instance into a private subnet. Associate an IPv6 address with the instance during launch. Configure a security group, a network ACL, and route tables for IPv6 communication. Create a route that directs all IPv6

traffic from the private subnet to an egress-only internet gateway.

ANSWER: C E

Explanation:

To test IPv6 connectivity for applications in a VPC, the instances must receive IPv6 addresses, the subnet and instance-level controls must allow IPv6 traffic, and the route tables must send IPv6 traffic to the correct gateway. Launching an Amazon EC2 instance into a public subnet, assigning an IPv6 address, and creating a route for IPv6 traffic to an internet gateway is correct for public IPv6 connectivity. In AWS, an internet gateway supports IPv6 internet access when the subnet route table includes an IPv6 default route such as `::/0` to the internet gateway, and security groups and network ACLs allow the required traffic.

Launching an Amazon EC2 instance into a private subnet, assigning an IPv6 address, and creating a route for IPv6 traffic to an egress-only internet gateway is also correct when the instance needs outbound IPv6 internet access without accepting unsolicited inbound IPv6 connections. An egress-only internet gateway is specifically designed for outbound-only IPv6 traffic from private subnets. These are the standard AWS patterns for IPv6-enabled public and private subnet connectivity. See [AWS IPv6 for VPCs](#) and [egress-only internet gateway documentation](#).

QUESTION NO: 102

A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services. Database credentials are hard-coded on each instance. SSH keys for command-line remote access are stored in a secured Amazon S3 bucket. The company has asked its solutions architect to improve the security posture of the architecture without adding operational complexity.

Which combination of steps should the solutions architect take to accomplish this? (Select THREE.)

- A. Use Amazon EC2 instance profiles with an IAM role.
- B. Use AWS Secrets Manager to store access keys and secret access keys.
- C. Use AWS Systems Manager Parameter Store to store database credentials.
- D. Use a secure fleet of Amazon EC2 bastion hosts (or remote access).
- E. Use AWS KMS to store database credentials.
- F. Use AWS Systems Manager Session Manager for remote access

ANSWER: A C F

Explanation:

Use Amazon EC2 instance profiles with an IAM role is correct because it removes the need to place long-term AWS access keys and secret access keys on instances. Applications running on the instance can retrieve temporary, automatically rotated credentials from the instance metadata service, which is the AWS-recommended approach for granting EC2 workloads access to AWS services. See [IAM roles for Amazon EC2](#).

Use AWS Systems Manager Parameter Store to store database credentials is also correct because it centralizes configuration and secrets outside the instance and can use encrypted SecureString parameters backed by AWS KMS. This avoids hard-coding credentials while keeping operations simple for applications that can retrieve parameters at runtime. See [AWS Systems Manager Parameter Store](#).

Use AWS Systems Manager Session Manager for remote access is correct because Session Manager provides managed shell access to instances without storing SSH keys, opening inbound SSH ports, or operating bastion hosts. It also integrates with IAM for access control and can log session activity, improving security without adding infrastructure to manage.

QUESTION NO: 103

A company stores a static website on Amazon S3. AWS Lambda functions retrieve content from an S3 bucket and serve the content as a website. An Application Load Balancer (ALB) directs incoming traffic to the Lambda functions. An Amazon CloudFront distribution routes requests to the ALB.

The company has set up an AWS Certificate Manager (ACM) certificate on the HTTPS listener of the ALB.

The company needs all users to communicate with the website through HTTPS. HTTP users must not receive an error.

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

- A. Configure the ALB with a TCP listener on port 443 for passthrough to backend systems.
- B. Create an S3 bucket policy that denies access to the S3 bucket if the `aws:SecureTransport` request is false.
- C. Configure HTTP to HTTPS redirection on the S3 bucket.
- D. Set the origin protocol policy to HTTPS Only for CloudFront.
- E. Set the viewer protocol policy to HTTPS Only for CloudFront.
- F. Set the viewer protocol policy to Redirect HTTP to HTTPS for CloudFront.

ANSWER: B D F

Explanation:

The correct combination is to enforce HTTPS across the full request path while redirecting users who arrive over HTTP. "Set the viewer protocol policy to Redirect HTTP to HTTPS for CloudFront." is required because CloudFront will return an HTTP redirect to the browser instead of an error, allowing HTTP users to reach the site securely. "Set the origin protocol policy to HTTPS Only for CloudFront." ensures that CloudFront communicates with the Application Load Balancer over HTTPS, which is supported because the ALB already has an ACM certificate on its HTTPS listener. "Create an S3 bucket policy that denies access to the S3 bucket if the `aws:SecureTransport` request is false." enforces encrypted transport for requests to the S3 bucket, helping ensure that content retrieval from S3 does not occur over plain HTTP. This aligns with AWS guidance for requiring SSL/TLS access to S3 by using the `aws:SecureTransport` condition key. See the CloudFront documentation for [requiring HTTPS between viewers and CloudFront](#) and the Amazon S3 documentation for [requiring requests to use SSL/TLS](#).

QUESTION NO: 104

A company runs an application in the cloud that consists of a database and a website. Users can post data to the website, have the data processed, and have the data sent back to them in an email. Data is stored in a MySQL database running on an Amazon EC2 instance. The database is running in a VPC with two private subnets. The website is running on Apache Tomcat in a single EC2 instance in a different VPC with one public subnet. There is a single VPC peering connection between the database and website VPC.

The website has suffered several outages during the last month due to high traffic.

Which actions should a solutions architect take to increase the reliability of the application? (Select THREE.)

- A. Place the Tomcat server in an Auto Scaling group with multiple EC2 instances behind an Application Load Balancer.
- B. Provision an additional VPC peering connection.
- C. Migrate the MySQL database to Amazon Aurora with one Aurora Replica.
- D. Provision two NAT gateways in the database VPC.
- E. Move the Tomcat server to the database VPC.
- F. Create an additional public subnet in a different Availability Zone in the website VPC.

ANSWER: A C F

Explanation:

Place the Tomcat server in an Auto Scaling group with multiple EC2 instances behind an Application Load Balancer is correct because it removes the single web server as a single point of failure and allows the website tier to scale out when traffic increases. An Application Load Balancer can distribute HTTP traffic across healthy Tomcat instances, while Amazon EC2 Auto Scaling can replace failed instances and add capacity based on demand. Create an additional public subnet in a different Availability Zone in the website VPC is also correct because a highly available load-balanced architecture should span multiple Availability Zones; the additional subnet enables the load balancer and Auto Scaling group to place resources across AZs for better fault tolerance. Migrate the MySQL database to Amazon Aurora with one Aurora Replica is correct because Aurora provides a managed, highly available MySQL-compatible database platform, and an Aurora Replica can improve read scalability and provide a failover target. Together, these changes improve reliability across both the web tier and the database tier. See AWS guidance for [using Elastic Load Balancing with Auto Scaling](#) and [Amazon Aurora high availability](#).

QUESTION NO: 105

A large company is running a popular web application. The application runs on several Amazon EC2 Linux instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the instances in the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager is configured, and AWS Systems Manager Agent is running on all the EC2 instances. The company recently released a new version of the application. Some EC2 instances are now being marked as unhealthy and are being terminated. As a result, the application is running at reduced capacity. A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive. How should the solutions architect gain access to an EC2 instance to troubleshoot the issue?

- A. Suspend the Auto Scaling group's HealthCheck scaling process. Use Session Manager to log in to an instance that is marked as unhealthy.
- B. Enable EC2 instance termination protection. Use Session Manager to log in to an instance that is marked as unhealthy.
- C. Set the termination policy to OldestInstance on the Auto Scaling group. Use Session Manager to log in to an instance that is marked an unhealthy.
- D. Suspend the Auto Scaling group's Terminate process. Use Session Manager to log in to an instance that is marked as unhealthy.

ANSWER: D

Explanation:

Suspend the Auto Scaling group's Terminate process. Use Session Manager to log in to an instance that is marked as unhealthy is the correct approach because it preserves an affected instance long enough for direct troubleshooting. When Amazon EC2 Auto Scaling determines that an instance is unhealthy, it normally terminates the instance and launches a replacement to maintain desired capacity. By suspending the Terminate process, the Auto Scaling group cannot terminate instances, including instances that have been marked unhealthy, which gives the solutions architect time to inspect the running environment, application files, local logs, dependencies, configuration, and runtime state before the instance is removed. AWS Systems Manager Session Manager is already configured and the SSM Agent is running, so the architect can connect securely to an instance in the private subnet without opening inbound SSH access, using a bastion host, or requiring a public IP address. This aligns with AWS guidance for temporarily suspending Auto Scaling processes during troubleshooting and using Session Manager for secure operational access to managed instances. See the AWS documentation for [suspending and resuming Auto Scaling processes](#) and [AWS Systems Manager Session Manager](#).

QUESTION NO: 106

A company is using AWS Organizations to manage multiple AWS accounts. For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts. A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of CloudFormation stacks. Trusted access has been enabled in Organizations. What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

- A. Create a stack set in the Organizations member accounts. Use service-managed permissions. Set deployment options to deploy to an organization. Use CloudFormation StackSets drift detection.

B. Create stacks in the Organizations member accounts. Use self-service permissions. Set deployment options to deploy to an organization. Enable the CloudFormation StackSets automatic deployment.

C. Create a stack set in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets automatic deployment.

D. Create stacks in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets drift detection.

ANSWER: C

Explanation:

Create a stack set in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets automatic deployment. is correct because AWS CloudFormation StackSets integrates with AWS Organizations through trusted access and service-managed permissions. With this model, the administrator creates the stack set from the organization management account, or from a registered delegated administrator, and CloudFormation handles the required execution roles in target member accounts. This is the intended approach for centrally deploying the same CloudFormation-defined resource, such as an Amazon SNS topic, across all accounts in an organization or selected organizational units. Setting the deployment target to the organization ensures the stack instances are created in the member accounts covered by the deployment scope. Enabling automatic deployment is important because it keeps the deployment aligned as the organization changes: when new accounts are added to the targeted organization or OU, StackSets can automatically create the required stack instances in those accounts. This matches the requirement to consistently deploy the SNS topic across all AWS accounts managed by AWS Organizations. See the AWS documentation for [enabling trusted access with AWS Organizations](#) and [automatic deployments for service-managed StackSets](#).

QUESTION NO: 107

A company wants to establish a dedicated connection between its on-premises infrastructure and AWS. The company is setting up a 1 Gbps AWS Direct Connect connection to its account VPC. The architecture includes a transit gateway and a Direct Connect gateway to connect multiple VPCs and the on-premises infrastructure. The company must connect to VPC resources over a transit VIF by using the Direct Connect connection. Which combination of steps will meet these requirements? (Choose two.)

A. Update the 1 Gbps Direct Connect connection to 10 Gbps.

B. Advertise the on-premises network prefixes over the transit VIF.

C. Advertise the VPC prefixes from the Direct Connect gateway to the on-premises network over the transit VIF.

D. Update the Direct Connect connection's MACsec encryption mode attribute to `must_encrypt`.

E. Associate a MACsec Connection Key Name/Connectivity Association Key (CKN/CAK) pair with the Direct Connect connection.

ANSWER: B C

Explanation:

To connect on-premises infrastructure to VPC resources through AWS Direct Connect by using a transit virtual interface, the routing must be exchanged across the transit VIF with the Direct Connect gateway and associated transit gateway.

“Advertise the on-premises network prefixes over the transit VIF” is correct because AWS needs to learn the customer network routes through BGP so traffic from VPCs attached to the transit gateway can return to the on-premises environment.

“Advertise the VPC prefixes from the Direct Connect gateway to the on-premises network over the transit VIF” is also correct because the on-premises routers need to learn the AWS-side reachable prefixes so they can route traffic to resources in the VPCs through the Direct Connect connection. This is the intended model for a transit virtual interface: it terminates on a Direct Connect gateway, which is associated with a transit gateway to provide scalable connectivity to multiple VPCs. AWS documentation describes transit virtual interfaces as the mechanism for accessing transit gateways through Direct Connect gateways, with route propagation handled through BGP advertisements. See [AWS Direct Connect virtual interfaces](#) and [Direct Connect gateways](#).

QUESTION NO: 108

A company is using an on-premises Active Directory service for user authentication. The company wants to use the same authentication service to sign in to the company 's AWS accounts, which are using AWS Organizations. AWS Site-to-Site VPN connectivity already exists between the on-premises environment and all the company 's AWS accounts.

The company 's security policy requires conditional access to the accounts based on user groups and roles. User identities must be managed in a single location.

Which solution will meet these requirements?

- A.** Configure AWS Single Sign-On (AWS SSO) to connect to Active Directory by using SAML 2.0. Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protocol. Grant access to the AWS accounts by using attribute-based access controls (ABACs).
- B.** Configure AWS Single Sign-On (AWS SSO) by using AWS SSO as an identity source. Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protocol. Grant access to the AWS accounts by using AWS SSO permission sets.
- C.** In one of the company 's AWS accounts, configure AWS Identity and Access Management (IAM) to use a SAML 2.0 identity provider. Provision IAM users that are mapped to the federated users. Grant access that corresponds to appropriate groups in Active Directory. Grant access to the required AWS accounts by using cross-account IAM users.
- D.** In one of the company 's AWS accounts, configure AWS Identity and Access Management (IAM) to use an OpenID Connect (OIDC) identity provider. Provision IAM roles that grant access to the AWS account for the federated users that correspond to appropriate groups in Active Directory. Grant access to the required AWS accounts by using cross-account IAM roles.
- E.** Configure AWS IAM Identity Center (AWS Single Sign-On) to use the on-premises Active Directory through AWS Directory Service (for example, AD Connector). Grant access to the AWS accounts in AWS Organizations by assigning Active Directory groups to IAM Identity Center permission sets.

ANSWER: E

Explanation:

Configure AWS IAM Identity Center (AWS Single Sign-On) to use the on-premises Active Directory through AWS Directory Service (for example, AD Connector). Grant access to the AWS accounts in AWS Organizations by assigning Active Directory groups to IAM Identity Center permission sets. is correct because IAM Identity Center is the AWS-recommended service for centralized workforce access across multiple AWS accounts in an organization. By connecting IAM Identity Center to Microsoft Active Directory through AWS Directory Service, users continue to authenticate with their existing corporate credentials, and identities remain managed in Active Directory as the single source of truth. The existing Site-to-Site VPN connectivity supports the network path needed for AWS Directory Service to communicate with the on-premises directory. Access can then be controlled by assigning Active Directory users or groups to AWS accounts and permission sets. IAM Identity Center permission sets are provisioned into target AWS accounts as IAM roles, which provides role-based access while keeping account access administration centralized. This directly supports group-based authorization, centralized identity management, and scalable multi-account access through AWS Organizations. See [AWS IAM Identity Center Active Directory identity source documentation](#) and [IAM Identity Center permission sets documentation](#).

QUESTION NO: 109

A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.

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

- A.** Use AWS CloudFormation StackSets to deploy AWS Config rules on production accounts.
- B.** Create a new AWS Control Tower landing zone in an existing developer account. Create OUs for accounts. Add production and development accounts to production and development OUs, respectively.

- C. Create a new AWS Control Tower landing zone in the company's management account. Add production and development accounts to production and development OUs, respectively.
- D. Invite existing accounts to join the organization in AWS Organizations. Create SCPs to ensure compliance.
- E. Create a guardrail from the management account to detect EBS encryption.
- F. Create a guardrail for the production OU to detect EBS encryption.

ANSWER: C D F

Explanation:

The correct approach is to use AWS Control Tower from the company's management account, organize accounts into separate production and development OUs, and apply the EBS encryption control only to the production OU. AWS Control Tower is designed to create a governed multi-account landing zone with built-in blueprints, account baselines, and controls, which matches the requirement for a standardized governance solution. Placing production and development accounts into separate OUs allows the company to scope controls so that only current and future production accounts are affected. Existing accounts can be invited into the AWS Organizations organization and then governed from the Control Tower environment. AWS Control Tower controls are applied at the OU level, so enabling the EBS encryption detection guardrail on the production OU ensures that accounts in that OU are continuously evaluated for compliance, including newly added production accounts. Preventive governance can also be implemented through service control policies as part of enforcing account-level compliance. For more information, see [AWS Control Tower](#) and [AWS Control Tower controls](#).

QUESTION NO: 110

A company is deploying a third-party firewall appliance solution from AWS Marketplace to monitor and protect traffic that leaves the company's AWS environments. The company wants to deploy this appliance into a shared services VPC and route all outbound internet-bound traffic through the appliances.

A solutions architect needs to recommend a deployment method that prioritizes reliability and minimizes failover time between firewall appliances within a single AWS Region. The company has set up routing from the shared services VPC to other VPCs.

Which steps should the solutions architect recommend to meet these requirements? (Select THREE)

- A. Deploy two firewall appliances into the shared services VPC, each in a separate Availability Zone.
- B. Create a new Network Load Balancer in the shared services VPC. Create a new target group, and attach it to the new Network Load Balancer. Add each of the firewall appliance instances to the target group.
- C. Create a new Gateway Load Balancer in the shared services VPC. Create a new target group, and attach it to the new Gateway Load Balancer. Add each of the firewall appliance instances to the target group.
- D. Create a VPC interface endpoint. Add a route to the route table in the shared services VPC. Designate the new endpoint as the next hop for traffic that enters the shared services VPC from other VPCs.
- E. Deploy two firewall appliances into the shared services VPC, each in the same Availability Zone.
- F. Create a VPC Gateway Load Balancer endpoint. Add a route to the route table in the shared services VPC. Designate the new endpoint as the next hop for traffic that enters the shared services VPC from other VPCs.

ANSWER: A C F

Explanation:

Deploying two firewall appliances into the shared services VPC, each in a separate Availability Zone, is correct because it provides Availability Zone-level resilience for the inspection layer. A Gateway Load Balancer is the AWS-recommended load balancing service for transparent insertion of third-party virtual appliances such as firewalls, intrusion detection systems, and deep packet inspection appliances. Creating a Gateway Load Balancer with a target group that contains the firewall appliance instances allows traffic flows to be distributed across healthy appliances while preserving flow symmetry, which is important for stateful inspection. Creating a VPC Gateway Load Balancer endpoint and using route tables to designate that

endpoint as the next hop is also correct because GWLB endpoints are used to steer traffic from VPC route tables to the Gateway Load Balancer service for inspection. This design minimizes failover time because the Gateway Load Balancer continuously monitors target health and forwards new flows only to healthy appliances, while the endpoint-based routing model supports centralized inspection through the shared services VPC. AWS documents this pattern for deploying virtual appliances with Gateway Load Balancer and routing traffic through Gateway Load Balancer endpoints. See [Gateway Load Balancer](#) and [Gateway Load Balancer endpoints](#).

QUESTION NO: 111

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 of "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

Explanation:

Activating the costCenter user-defined tag in the management account and using AWS Cost and Usage Reports is the most accurate approach for allocating these security-tool costs across an AWS Organizations environment. In consolidated billing, the management account is the appropriate place to activate user-defined cost allocation tags so that tagged usage from member accounts can be included in billing and cost reporting. After activation, the tag becomes available as a cost allocation dimension, allowing costs for resources tagged with costCenter=compliance to be separated from other spend.

AWS Cost and Usage Reports provide the most detailed billing data available from AWS, including line-item usage, costs, account information, service details, and activated cost allocation tags. Storing the report in an Amazon S3 bucket in the management account allows the company to centrally analyze costs for tagged EC2-based security tooling across all accounts and charge those costs back to the compliance team's account with the highest practical accuracy. AWS documents that cost allocation tags can be activated for billing analysis and that Cost and Usage Reports contain comprehensive cost and usage data: [AWS cost allocation tags](#) and [AWS Cost and Usage Reports](#).

QUESTION NO: 112

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 D

Explanation:

The correct solution is to use VPC sharing with AWS Resource Access Manager. "Enable resource sharing from the AWS Organizations management account" is required because sharing resources with accounts or organizational units in AWS Organizations must first be enabled from the organization's management account. After this integration is enabled, resources can be shared broadly with member accounts or OUs without needing individual invitations. "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." is correct because VPC sharing allows the VPC owner account to centrally own and manage the VPC, route tables, network ACLs, and subnets, while participant accounts can launch supported AWS resources into the shared subnets. This matches the requirement that the infrastructure team manages the network while individual accounts can create resources within subnets. AWS documents this pattern as VPC sharing using AWS RAM, where subnets are shared with other accounts, OUs, or the organization. References: [AWS VPC sharing](#) and [AWS RAM resource sharing](#).

QUESTION NO: 113

A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application's data. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record. The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an active-active strategy. What should a solutions architect recommend to meet these requirements?

- A.** Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region. Configure the CloudWatch alarm to invoke the Lambda function.
- B.** Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealthy. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.
- C.** Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Region. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs. Remove the read replica. Replace the read replica with a standalone RDS DB instance. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.
- D.** Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted targets. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region. Configure the CloudWatch alarm to invoke the Lambda function.

ANSWER: B

Explanation:

Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealthy. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs. is the correct recommendation because it implements an automated pilot-light style disaster recovery pattern. Route 53 DNS failover can monitor the primary application endpoint with a health check and automatically return the secondary ALB record when the primary health check fails. That provides automatic client redirection without requiring an active-active deployment. The SNS-to-Lambda workflow supplies the missing recovery automation in the backup Region: the Lambda function can promote the cross-Region RDS read replica to a standalone writer database and update the Auto Scaling group capacity from zero so application instances launch behind the standby ALB. This keeps steady-state costs low while still enabling recovery within a short RTO target, assuming appropriate DNS TTLs and preconfigured application infrastructure. AWS documents Route 53 DNS failover with health checks in the [Route 53 Developer Guide](#), and RDS read replica promotion in the [Amazon RDS User Guide](#).

QUESTION NO: 114

A company has a web application that securely uploads pictures and videos to an Amazon S3 bucket. The company requires that only authenticated users are allowed to post content. The application generates a presigned URL that is used to upload objects through a browser interface. Most users are reporting slow upload times for objects larger than 100 MB.

What can a Solutions Architect do to improve the performance of these uploads while ensuring only authenticated users are allowed to post content?

- A.** Set up an Amazon API Gateway with an edge-optimized API endpoint that has a resource as an S3 service proxy. Configure the PUT method for this resource to expose the S3 PutObject operation. Secure the API Gateway using a COGNITO_USER_POOLS authorizer. Have the browser interface use API Gateway instead of the presigned URL to upload objects.
- B.** Set up an Amazon API Gateway with a regional API endpoint that has a resource as an S3 service proxy. Configure the PUT method for this resource to expose the S3 PutObject operation. Secure the API Gateway using an AWS Lambda authorizer. Have the browser interface use API Gateway instead of the presigned URL to upload API objects.
- C.** Enable an S3 Transfer Acceleration endpoint on the S3 bucket. Use the endpoint when generating the presigned URL. Have the browser interface upload the objects to this URL using the S3 multipart upload API.
- D.** Configure an Amazon CloudFront distribution for the destination S3 bucket. Enable PUT and POST methods for the CloudFront cache behavior. Update the CloudFront origin to use an origin access identity (OAI). Give the OAI user s3:PutObject permissions in the bucket policy. Have the browser interface upload objects using the CloudFront distribution

ANSWER: C

Explanation:

Enable an S3 Transfer Acceleration endpoint on the S3 bucket. Use the endpoint when generating the presigned URL. Have the browser interface upload the objects to this URL using the S3 multipart upload API. is correct because it combines two AWS-recommended mechanisms for improving large uploads to Amazon S3 while preserving authenticated access control. S3 Transfer Acceleration routes uploads through the nearest AWS edge location and then across the optimized AWS global network to the destination bucket, which can significantly improve performance for users who are geographically distant from the bucket's Region. AWS explicitly supports using Transfer Acceleration with presigned URLs, so the application can continue generating upload URLs only after users authenticate successfully. For objects larger than 100 MB, the multipart upload API is also a best practice because it splits the object into parts that can be uploaded in parallel and retried independently if a part fails. Together, Transfer Acceleration and multipart upload address latency, throughput, and reliability for large browser-based uploads without requiring the application to proxy the file data. See [Amazon S3 Transfer Acceleration](#) and [Amazon S3 multipart upload overview](#).

QUESTION NO: 115

A company is developing a gene reporting device that will collect genomic information to assist researchers with collecting large samples of data from a diverse population. The device will push 8 KB of genomic data every second to a data platform that will need to process and analyze the data and provide information back to researchers. The data platform must meet the following requirements: • Provide near-real-time analytics of the inbound genomic data • Ensure the data is flexible, parallel, and durable • Deliver results of processing to a data warehouse Which strategy should a solutions architect use to meet these requirements?

- A.** Use Amazon Kinesis Data Firehose to collect the inbound sensor data, analyze the data with Kinesis clients, and save the results to an Amazon RDS instance.
- B.** Use Amazon Kinesis Data Streams to collect the inbound sensor data, analyze the data with Kinesis clients, and save the results to an Amazon Redshift cluster using Amazon EMR.
- C.** Use Amazon S3 to collect the inbound device data, analyze the data from Amazon SQS with Kinesis, and save the results to an Amazon Redshift cluster.
- D.** Use an Amazon API Gateway to put requests into an Amazon SQS queue, analyze the data with an AWS Lambda function, and save the results to an Amazon Redshift cluster using Amazon EMR.

ANSWER: B

Explanation:

Use Amazon Kinesis Data Streams to collect the inbound sensor data, analyze the data with Kinesis clients, and save the results to an Amazon Redshift cluster using Amazon EMR is the correct strategy. Amazon Kinesis Data Streams is designed for high-throughput, near-real-time ingestion of streaming records from many producers, which fits a device continuously sending genomic payloads. Its shard-based model provides parallelism and scalability, while records are replicated across Availability Zones for durability during the configured retention period. Kinesis client applications can consume the stream concurrently and process records with low latency, enabling near-real-time analytics over the inbound genomic data.

Amazon EMR is appropriate for running distributed processing frameworks, such as Apache Spark, against streaming data and preparing aggregated or transformed results. Amazon Redshift is AWS's managed data warehouse service, so storing the processed output in Amazon Redshift aligns with the requirement to deliver results to a data warehouse for researchers and downstream analytics. See the AWS documentation for [Amazon Kinesis Data Streams](#) and [Amazon EMR](#).

QUESTION NO: 116

A company is planning to set up a REST API application on AWS. The application team wants to set up a new identity store on AWS. The IT team does not want to maintain any infrastructure or servers for this deployment.

What is the MOST operationally efficient solution that meets these requirements?

- A.** Deploy the application as AWS Lambda functions. Set up Amazon API Gateway REST API endpoints for the application. Create a Lambda function, and configure a Lambda authorizer.
- B.** Deploy the application in AWS AppSync, and configure AWS Lambda resolvers. Set up an Amazon Cognito user pool, and configure AWS AppSync to use the user pool for authorization.
- C.** Deploy the application as AWS Lambda functions. Set up Amazon API Gateway REST API endpoints for the application. Set up an Amazon Cognito user pool, and configure an Amazon Cognito authorizer.
- D.** Deploy the application in Amazon Elastic Kubernetes Service (Amazon EKS) clusters. Set up an Application Load Balancer for the EKS pods. Set up an Amazon Cognito user pool and service pod for authentication.

ANSWER: C

Explanation:

Deploying the application as AWS Lambda functions with Amazon API Gateway REST API endpoints and using an Amazon Cognito user pool with an Amazon Cognito authorizer is the most operationally efficient solution. This design is fully serverless: Lambda runs application code without requiring the team to provision or manage compute infrastructure, and API Gateway provides managed REST API hosting, routing, scaling, throttling, and integration with Lambda. Amazon Cognito

user pools provide a managed identity store for application users, including sign-up, sign-in, user directories, password policies, token issuance, and federation capabilities without requiring the IT team to operate authentication servers. API Gateway REST APIs can use a Cognito user pool authorizer to authenticate requests by validating tokens issued by Cognito before invoking the backend Lambda functions. This directly matches the requirements for a REST API, a new AWS-native identity store, and no server or infrastructure maintenance. See the AWS documentation for [using Amazon Cognito user pools as authorizers in API Gateway](#) and [Amazon Cognito user pools](#).

QUESTION NO: 117

A company has developed a new release of a popular video game and wants to make it available for public download. The new release package is approximately 5 GB in size. The company provides downloads for existing releases from a Linux-based, publicly facing FTP site hosted in an on-premises data center. The company expects the new release will be downloaded by users worldwide. The company wants a solution that provides improved download performance and low transfer costs, regardless of a user's location.

Which solutions will meet these requirements?

- A.** Store the game files on Amazon EBS volumes mounted on Amazon EC2 instances within an Auto Scaling group. Configure an FTP service on the EC2 instances. Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to download the package.
- B.** Store the game files on Amazon EFS volumes that are attached to Amazon EC2 instances within an Auto Scaling group. Configure an FTP service on each of the EC2 instances. Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to download the package.
- C.** Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Use Amazon CloudFront for the website. Publish the game download URL for users to download the package.
- D.** Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Set Requester Pays for the S3 bucket. Publish the game download URL for users to download the package.

ANSWER: C

Explanation:

Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Use Amazon CloudFront for the website. Publish the game download URL for users to download the package. is correct because it uses Amazon S3 as a highly durable, scalable origin for the 5 GB game package and Amazon CloudFront as a global content delivery network. CloudFront caches content at edge locations and regional edge caches around the world, so users download the file from a location closer to them instead of repeatedly pulling it from a single data center or Region. This directly improves global download performance and reduces latency for a worldwide user base. CloudFront is also designed for high-volume content distribution and can reduce origin load and data transfer costs by serving cached objects from edge locations. Route 53 can provide a friendly DNS name for the download endpoint, while S3 static website hosting or an S3 origin can host the downloadable object reliably without requiring the company to manage FTP servers or compute fleets. This architecture is a common AWS pattern for public software distribution at scale. See the AWS documentation for [Amazon CloudFront](#) and [hosting a static website with Amazon S3](#).

QUESTION NO: 118

A company uses a load balancer to distribute traffic to Amazon EC2 instances in a single Availability Zone. The company is concerned about security and wants a solutions architect to re-architect the solution to meet the following requirements:

- Inbound requests must be filtered for common vulnerability attacks.
- Rejected requests must be sent to a third-party auditing application.
- All resources should be highly available.

Which solution meets these requirements?

- A.** Configure a Multi-AZ Auto Scaling group using the application's AMI. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Use Amazon Inspector to monitor traffic to the ALB and EC2 instances. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application.

B. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Logs. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.

C. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as targets. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

D. Configure a Multi-AZ Auto Scaling group using the application's AMI. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the WebACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

ANSWER: D

Explanation:

The solution that configures a Multi-AZ Auto Scaling group, places it behind an Application Load Balancer, associates AWS WAF with the load balancer, and sends AWS WAF logs to Amazon Kinesis Data Firehose best satisfies the requirements. A Multi-AZ Auto Scaling group provides highly available application capacity across Availability Zones, while the Application Load Balancer distributes traffic to healthy targets. AWS WAF can be associated with an Application Load Balancer to inspect HTTP/S requests and block requests that match web ACL rules, including managed rule groups that protect against common web exploits such as SQL injection, cross-site scripting, and known bad inputs. AWS WAF logging can deliver detailed request logs to Kinesis Data Firehose, which can then forward the data to supported destinations, including third-party HTTP endpoints or downstream auditing systems. This allows blocked or rejected request activity to be exported for external audit processing while keeping the application tier highly available. See the AWS documentation for [AWS WAF logging](#) and [AWS Managed Rules rule groups](#).

QUESTION NO: 119

A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC.

A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.

Which solution meets these requirements?

A. Provision a Direct Connect gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the Direct Connect gateway. Connect the Direct Connect gateway to the single VPC

B. Keep the existing private virtual interface. Create the second Direct Connect connection. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC

C. Keep the existing private virtual interface. Create the second Direct Connect connection. Create a new public virtual interface on the new connection, and connect the new public virtual interface to the single VPC

D. Provision a transit gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gateway. Associate the transit gateway with the single VPC

ANSWER: A

Explanation:

Provision a Direct Connect gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection. Create a new private virtual interface on each connection, and connect both private

virtual interfaces to the Direct Connect gateway. Connect the Direct Connect gateway to the single VP is the correct solution because an AWS Direct Connect gateway is designed to let private virtual interfaces on Direct Connect connections reach VPCs across multiple AWS Regions. By terminating a private virtual interface from each Direct Connect connection on the same Direct Connect gateway, the company gets redundant physical connectivity and redundant BGP paths from the on-premises network into AWS. The Direct Connect gateway can then be associated with the current VPC's virtual private gateway and later associated with additional VPCs in other supported Regions, allowing the same pair of Direct Connect connections to serve the company's regional expansion without provisioning separate Direct Connect circuits per Region. AWS documents Direct Connect gateways as a way to connect Direct Connect connections to VPCs in any AWS Region, and private virtual interfaces are the appropriate interface type for private VPC connectivity. See the AWS Direct Connect gateway documentation at [AWS Direct Connect gateways](#) and private virtual interface guidance at [Working with virtual interfaces](#).

QUESTION NO: 120

A company runs applications in hundreds of production AWS accounts. The company uses AWS Organizations with all features enabled and has a centralized backup operation that uses AWS Backup. The company is concerned about ransomware attacks. To address this concern, the company has created a new policy that all backups must be resilient to breaches of privileged-user credentials in any production account. Which combination of steps will meet this new requirement? (Choose three.)

- A. Implement cross-account backup with AWS Backup vaults in designated non-production accounts.
- B. Add an SCP that restricts the modification of AWS Backup vaults.
- C. Implement AWS Backup Vault Lock in compliance mode.
Implement least privilege access for the IAM service role that is assigned to AWS Backup.
- D. Configure the backup frequency, lifecycle, and retention period to ensure that at least one backup always exists in the cold tier.
- E. Configure AWS Backup to write all backups to an Amazon S3 bucket in a designated non- production account. Ensure that the S3 bucket has S3 Object Lock enabled.

ANSWER: A B C

Explanation:

The best combination is to use cross-account backup vaults in designated non-production accounts, restrict backup-vault modification with service control policies, and enable AWS Backup Vault Lock in compliance mode while applying least privilege to the AWS Backup IAM service role. Cross-account backup creates backup copies outside the production accounts, so a compromised privileged identity in a production account cannot directly delete or tamper with the protected copies in the backup account. AWS Organizations service control policies add preventive guardrails across member accounts, including controls that can deny actions that modify or delete backup vaults and recovery points, even for highly privileged principals in those accounts. AWS Backup Vault Lock in compliance mode provides write-once-read-many protection for recovery points by enforcing retention settings that cannot be shortened or removed after the cooling-off period, helping protect backups from malicious or accidental deletion. Least privilege for the AWS Backup service role further reduces the blast radius by granting only the permissions required for backup and restore operations. Together, these controls align with AWS guidance for ransomware-resilient backup architectures. See [AWS Backup cross-account backup](#) and [AWS Backup Vault Lock](#).

QUESTION NO: 121

A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.

What should the solutions architect do to meet these requirements?

- A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows File Server Use the SMB share to host the VMware data store. Use VM Import/Export to move the VMs to Amazon EC2.

B. Use the VMware vSphere client to export the application as an image in Open Virtualization Format (OVF) format. Create an Amazon S3 bucket to store the image in the destination AWS Region. Create and apply an IAM role for VM Import. Use the AWS CLI to run the EC2 import command.

C. Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFS) share. Create a backup copy to the shared folder. Sign in to the AWS Management Console and create an AMI from the backup copy. Launch an EC2 instance that is based on the AMI.

D. Create a managed-instance activation for a hybrid environment in AWS Systems Manager. Download and install Systems Manager Agent on the on-premises VM. Register the VM with Systems Manager to be a managed instance. Use AWS Backup to create a snapshot of the VM and create an AMI. Launch an EC2 instance that is based on the AMI.

ANSWER: B

Explanation:

Use the VMware vSphere client to export the application as an image in Open Virtualization Format (OVF) format. Create an Amazon S3 bucket to store the image in the destination AWS Region. Create and apply an IAM role for VM Import. Use the AWS CLI to run the EC2 import command. This is correct because VM Import/Export is the AWS-supported mechanism for importing existing virtual machine images from on-premises virtualization environments, including VMware, into Amazon EC2 while preserving the installed operating system, applications, and configuration. The typical workflow is to export the VM from VMware as an OVF/OVA or supported disk image, upload the image files to an Amazon S3 bucket in the target Region, configure the required vmimport IAM role and permissions, and then use the AWS CLI import-image or import-instance process to create an Amazon Machine Image or EC2 instance. This approach directly addresses the requirement to migrate the application as-is rather than rebuilding it or copying only application data. AWS documents this process under VM Import/Export and specifically supports importing VMs from virtualization platforms such as VMware. See the AWS VM Import/Export documentation for image import requirements and workflow: [Importing a VM as an image using VM Import/Export](#) and [What is VM Import/Export?](#).

QUESTION NO: 122

A company wants to run a custom network analysis software package to inspect traffic as traffic leaves and enters a VPC. The company has deployed the solution by using AWS CloudFormation on three Amazon EC2 instances in an Auto Scaling group. All network routing has been established to direct traffic to the EC2 instances. Whenever the analysis software stops working, the Auto Scaling group replaces an instance. The network routes are not updated when the instance replacement occurs. Which combination of steps will resolve this issue? (Choose three.)

A. Create alarms based on EC2 status check metrics that will cause the Auto Scaling group to replace the failed instance.

B. Update the CloudFormation template to install the Amazon CloudWatch agent on the EC2 instances. Configure the CloudWatch agent to send process metrics for the application.

C. Update the CloudFormation template to install AWS Systems Manager Agent on the EC2 instances. Configure Systems Manager Agent to send process metrics for the application.

D. Create an alarm for the custom metric in Amazon CloudWatch for the failure scenarios. Configure the alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

E. Create an AWS Lambda function that responds to the Amazon Simple Notification Service (Amazon SNS) message to take the instance out of service. Update the network routes to point to the replacement instance.

F. In the CloudFormation template, write a condition that updates the network routes when a replacement instance is launched.

ANSWER: B D E

Explanation:

The correct approach is to monitor the actual analysis software process, raise an actionable event when that process fails, and automate the routing update during replacement. Updating the CloudFormation template to install the Amazon CloudWatch agent on the EC2 instances is correct because the CloudWatch agent can collect process-level metrics, such as whether the custom network analysis process is running, by using procstat-style monitoring. Creating an alarm for the

custom metric in Amazon CloudWatch for the failure scenarios and publishing to an Amazon SNS topic is correct because it turns the process failure signal into an event that can trigger remediation. Creating an AWS Lambda function that responds to the Amazon SNS message to take the instance out of service and update the network routes to point to the replacement instance is also correct because VPC route entries that target an EC2 instance or network interface do not automatically follow Auto Scaling replacements; the route target must be programmatically changed, for example by using the EC2 [ReplaceRoute API](#). See AWS documentation for [collecting process metrics with the CloudWatch agent](#) and the [Amazon EC2 ReplaceRoute API](#).

QUESTION NO: 123

A company has an on-premises monitoring solution using a PostgreSQL database for persistence of events. The database is unable to scale due to heavy ingestion and it frequently runs out of storage.

The company wants to create a hybrid solution and has already set up a VPN connection between its network and AWS. The solution should include the following attributes:

- Managed AWS services to minimize operational complexity
- A buffer that automatically scales to match the throughput of data and requires no on-going administration.
- A visualization tool to create dashboards to observe events in near-real time.
- Support for semi-structured JSON data and dynamic schemas.

Which combination of components will enable the company to create a monitoring solution that will satisfy these requirements? (Select TWO.)

- A. Use Amazon Kinesis Data Firehose to buffer events. Create an AWS Lambda function to process and transform events.
- B. Create an Amazon Kinesis data stream to buffer events. Create an AWS Lambda function to process and transform events.
- C. Configure an Amazon Aurora PostgreSQL DB cluster to receive events. Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards.
- D. Configure Amazon Elasticsearch Service (Amazon ES) to receive events. Use the Kibana endpoint deployed with Amazon ES to create near-real-time visualizations and dashboards.
- E. Configure an Amazon Neptune DB instance to receive events. Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards.

ANSWER: A D

Explanation:

The correct combination is to use Amazon Kinesis Data Firehose to buffer events. Create an AWS Lambda function to process and transform events and configure Amazon Elasticsearch Service (Amazon ES) to receive events. Use the Kibana endpoint deployed with Amazon ES to create near-real-time visualizations and dashboards. Kinesis Data Firehose is a fully managed ingestion and delivery service that scales automatically and requires little operational administration, which fits the need for a managed buffer for high-throughput event data. It can invoke Lambda for transformation before delivering records to destinations such as Amazon OpenSearch Service, formerly Amazon Elasticsearch Service. See the AWS documentation for Firehose delivery streams and transformation support: [Amazon Data Firehose](#).

Amazon Elasticsearch Service, now Amazon OpenSearch Service, is designed for search, log analytics, and observability use cases. It stores and indexes JSON documents, supports flexible and dynamic schemas through mappings, and provides a dashboarding interface through Kibana/OpenSearch Dashboards for near-real-time visualization of ingested events. This aligns directly with the requirement to analyze semi-structured monitoring events and build operational dashboards. See [Amazon OpenSearch Service](#).

QUESTION NO: 124

A company wants to migrate its on-premises application to AWS. The database for the application stores structured product data and temporary user session data. The company needs to decouple the product data from the user session data. The

company also needs to implement replication in another AWS Region for disaster recovery. Which solution will meet these requirements with the HIGHEST performance?

- A.** Create an Amazon RDS DB instance with separate schemas to host the product data and the user session data. Configure a read replica for the DB instance in another Region.
- B.** Create an Amazon RDS DB instance to host the product data. Configure a read replica for the DB instance in another Region. Create a global datastore in Amazon ElastiCache for Memcached to host the user session data.
- C.** Create two Amazon DynamoDB global tables. Use one global table to host the product data. Use the other global table to host the user session data. Use DynamoDB Accelerator (DAX) for caching.
- D.** Create an Amazon RDS DB instance to host the product data. Configure a read replica for the DB instance in another Region. Create an Amazon DynamoDB global table to host the user session data.

ANSWER: D

Explanation:

Create an Amazon RDS DB instance to host the product data. Configure a read replica for the DB instance in another Region. Create an Amazon DynamoDB global table to host the user session data is correct because it separates the two data domains while using services that fit their access patterns. Amazon RDS is appropriate for structured product data that commonly requires relational querying, transactions, and managed database capabilities. A cross-Region RDS read replica provides an asynchronous copy in another AWS Region, which can be promoted during a disaster recovery event, supporting regional resilience for the product catalog database. See the AWS documentation for [Amazon RDS cross-Region read replicas](#).

Amazon DynamoDB global tables are a strong fit for temporary user session data because they provide low-latency, highly scalable key-value access and automatically replicate table data across selected AWS Regions. This enables session data to remain decoupled from the relational product database and available in another Region for disaster recovery or active-active application designs. DynamoDB is designed for single-digit millisecond performance at scale, and global tables extend that model across Regions. See [DynamoDB global tables](#) for details.

QUESTION NO: 125

A company is serving files to its customers through an SFTP server that is accessible over the internet. The SFTP server is running on a single Amazon EC2 instance with an Elastic IP address attached. Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication. The EC2 instance also has an attached security group that allows access from all customer IP addresses. A solutions architect must implement a solution to improve availability, minimize the complexity of infrastructure management, and minimize the disruption to customers who access files. The solution must not change the way customers connect. Which solution will meet these requirements?

- A.** Disassociate the Elastic IP address from the EC2 instance. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family server. Configure the Transfer Family server with a publicly accessible endpoint. Associate the SFTP Elastic IP address with the new endpoint. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.
- B.** Disassociate the Elastic IP address from the EC2 instance. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family server. Configure the Transfer Family server with a VPC-hosted, internet-facing endpoint. Associate the SFTP Elastic IP address with the new endpoint. Attach the security group with customer IP addresses to the new endpoint. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.
- C.** Disassociate the Elastic IP address from the EC2 instance. Create a new Amazon Elastic File System (Amazon EFS) file system to be used for SFTP file hosting. Create an AWS Fargate task definition to run an SFTP server. Specify the EFS file system as a mount in the task definition. Create a Fargate service by using the task definition, and place a Network Load Balancer (NLB) in front of the service. When configuring the service, attach the security group with customer IP addresses to the tasks that run the SFTP server. Associate the Elastic IP address with the NLB. Sync all files from the SFTP server to the S3 bucket.
- D.** Disassociate the Elastic IP address from the EC2 instance. Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hosting. Create a Network Load Balancer (NLB) with the Elastic IP address attached. Create an Auto Scaling group with EC2 instances that run an SFTP server. Define in the Auto Scaling group that instances

that are launched should attach the new multi-attach EBS volume. Configure the Auto Scaling group to automatically add instances behind the NLB. Configure the Auto Scaling group to use the security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launches. Sync all files from the SFTP server to the new multi-attach EBS volume.

ANSWER: B

Explanation:

Using AWS Transfer Family with a VPC-hosted, internet-facing endpoint, associating the existing Elastic IP address with that endpoint, and backing the server with Amazon S3 is the correct solution. AWS Transfer Family is a fully managed service for SFTP, so it removes the operational burden of maintaining an EC2-based SFTP server while improving availability through AWS-managed infrastructure. A VPC-hosted internet-facing endpoint supports Elastic IP addresses, which allows the company to preserve the same public IP address that customers already use. That directly satisfies the requirement to avoid changing the customer connection method. The endpoint can also use security groups, allowing the existing customer IP allow list behavior to be maintained. Storing files in Amazon S3 provides durable, highly available storage and integrates natively with Transfer Family as a backend. This approach gives customers the same SFTP protocol and SSH-based access pattern while modernizing the implementation to a managed, more resilient architecture. See the AWS documentation for [creating an AWS Transfer Family server in a VPC](#) and [AWS Transfer Family managed file transfer](#).

QUESTION NO: 126

A company has multiple AWS accounts as part of an organization created with AWS Organizations. Each account has a VPC in the us-east-2 Region and is used for either production or development workloads. Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other.

To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another.

Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

- A.** Modify the security groups assigned to development EC2 instances to block traffic from production EC2 instances. Modify the security groups assigned to production EC2 instances to block traffic from development EC2 instances.
- B.** Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached. Using the Network Manager feature of AWS Transit Gateway, create policies that restrict traffic between VPCs based on the value of this tag.
- C.** Create separate route tables for production and development traffic. Delete each account's association and route propagation to the default AWS Transit Gateway route table. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment.
- D.** Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached. Modify the AWS Transit Gateway routing table to route production tagged attachments to one another and development tagged attachments to one another.

ANSWER: C

Explanation:

Create separate route tables for production and development traffic. Delete each account's association and route propagation to the default AWS Transit Gateway route table. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment is correct because AWS Transit Gateway provides isolation through separate transit gateway route tables. Each VPC attachment can be associated with a specific transit gateway route table, and routes from attachments can be

propagated only into the route tables where connectivity is intended. By placing production VPC attachments in a production transit gateway route table and development VPC attachments in a development transit gateway route table, production VPCs learn routes only to other production VPCs, while development VPCs learn routes only to other development VPCs. Removing the default route table association and propagation is important because the default route table can otherwise unintentionally provide shared connectivity between all attachments. This design uses native Transit Gateway segmentation and routing controls, which is the appropriate way to enforce network-level isolation at scale across multiple accounts and VPCs. AWS documents that transit gateway route tables control packet routing between attachments and that attachments can be associated with and propagate routes to route tables. See [Transit gateway route tables](#) and [How transit gateways work](#).

QUESTION NO: 127

A company has implemented a new security requirement. According to the new requirement, the company must scan all traffic from corporate AWS instances in the company's VPC for violations of the company's security policies. As a result of these scans, the company can block access to and from specific IP addresses. To meet the new requirement, the company deploys a set of Amazon EC2 instances in private subnets to serve as transparent proxies. The company installs approved proxy server software on these EC2 instances. The company modifies the route tables on all subnets to use the corresponding EC2 instances with proxy software as the default route. The company also creates security groups that are compliant with the security policies and assigns these security groups to the EC2 instances. Despite these configurations, the traffic of the EC2 instances in their private subnets is not being properly forwarded to the internet. What should a solutions architect do to resolve this issue?

- A. Disable source/destination checks on the EC2 instances that run the proxy software.
- B. Add a rule to the security group that is assigned to the proxy EC2 instances to allow all traffic between instances that have this security group. Assign this security group to all EC2 instances in the VPC.
- C. Change the VPCs DHCP options set. Set the DNS server options to point to the addresses of the proxy EC2 instances.
- D. Assign one additional elastic network interface to each proxy EC2 instance. Ensure that one of these network interfaces has a route to the private subnets. Ensure that the other network interface has a route to the internet.

ANSWER: A

Explanation:

Disable source/destination checks on the EC2 instances that run the proxy software is correct because these instances are being used as forwarding devices rather than as normal endpoint instances. By default, Amazon EC2 performs a source/destination check on each instance, which means the instance must be either the source or the destination of any traffic that it receives. A transparent proxy or NAT-style inspection instance receives packets from other instances and forwards them onward to another destination, such as the internet. In that forwarding path, the proxy instance is not the original source or final destination for the traffic, so the default source/destination check can prevent the traffic from being forwarded correctly.

AWS specifically requires disabling source/destination checks for EC2 instances that perform network address translation, routing, or other traffic-forwarding functions. Since the route tables already direct default traffic to the proxy EC2 instances and the proxy software and security groups have been configured, disabling this EC2 networking safeguard is the necessary AWS-side change to allow the instances to operate as transparent network appliances. See the AWS documentation on [changing the source/destination check](#) and the NAT instance guidance in the [Amazon VPC User Guide](#).

QUESTION NO: 128

A company manages multiple AWS accounts by using AWS Organizations. Under the root OU, the company has two OUs: Research and DataOps. Because of regulatory requirements, all resources that the company deploys in the organization must reside in the ap-northeast-1 Region. Additionally, EC2 instances that the company deploys in the DataOps OU must use a predefined list of instance types. A solutions architect must implement a solution that applies these restrictions. The solution must maximize operational efficiency and must minimize ongoing maintenance. Which combination of steps will meet these requirements? (Choose two.)

- A. Create an IAM role in one account under the DataOps OU. Use the ec2:InstanceType condition key in an inline policy on the role to restrict access to specific instance type.

- B.** Create an IAM user in all accounts under the root OU. Use the `aws:RequestedRegion` condition key in an inline policy on each user to restrict access to all AWS Regions except `ap-northeast-1`.
- C.** Create an SCP. Use the `aws:RequestedRegion` condition key to restrict access to all AWS Regions except `ap-northeast-1`. Apply the SCP to the root OU.
- D.** Create an SCP. Use the `ec2:Region` condition key to restrict access to all AWS Regions except `ap-northeast-1`. Apply the SCP to the root OU, the DataOps OU, and the Research OU.
- E.** Create an SCP. Use the `ec2:InstanceType` condition key to restrict access to specific instance types. Apply the SCP to the DataOps OU.

ANSWER: C E

Explanation:

Create an SCP. Use the `aws:RequestedRegion` condition key to restrict access to all AWS Regions except `ap-northeast-1`. Apply the SCP to the root OU. is correct because a service control policy applied at the root of an AWS Organizations hierarchy establishes a centrally managed permissions guardrail for every account in the organization, including accounts in both the Research and DataOps OUs. The global condition key `aws:RequestedRegion` is the standard way to restrict API requests to an allowed Region, and applying it at the root minimizes ongoing administration as accounts are added or moved within the organization. See AWS guidance on Region restriction SCPs in the [AWS Organizations SCP examples](#).

Create an SCP. Use the `ec2:InstanceType` condition key to restrict access to specific instance types. Apply the SCP to the DataOps OU. is also correct because the instance type requirement applies only to EC2 usage within the DataOps OU. Using an SCP at that OU level enforces the allowed EC2 instance types across all current and future accounts in that OU without requiring per-account IAM policy maintenance. The `ec2:InstanceType` condition key is supported for controlling EC2 actions such as instance launches based on instance type, as documented in the [Amazon EC2 service authorization reference](#).

QUESTION NO: 129

A company operates a fleet of servers on premises and operates a fleet of Amazon EC2 instances in its organization in AWS Organizations. The company's AWS accounts contain hundreds of VPCs. The company wants to connect its AWS accounts to its on-premises network. AWS Site-to-Site VPN connections are already established to a single AWS account. The company wants to control which VPCs can communicate with other VPCs. Which combination of steps will achieve this level of control with the LEAST operational effort? (Choose three.)

- A.** Create a transit gateway in an AWS account. Share the transit gateway across accounts by using AWS Resource Access Manager (AWS RAM).
- B.** Configure attachments to all VPCs and VPNs.
- C.** Setup transit gateway route tables. Associate the VPCs and VPNs with the route tables.
- D.** Configure VPC peering between the VPCs.
- E.** Configure attachments between the VPCs and VPNs.
- F.** Setup route tables on the VPCs and VPNs.

ANSWER: A B C

Explanation:

The least operational effort for hundreds of VPCs across AWS accounts is to use AWS Transit Gateway as a central hub. "Create a transit gateway in an AWS account. Share the transit gateway across accounts by using AWS Resource Access Manager (AWS RAM)." is correct because AWS RAM lets an organization centrally share a transit gateway with other accounts, avoiding the need to build and manage large numbers of individual network relationships. "Configure attachments to all VPCs and VPNs." is correct because each participating VPC and the existing Site-to-Site VPN connectivity must be attached to the transit gateway so traffic can flow through the hub. "Setup transit gateway route tables. Associate the VPCs and VPNs with the route tables." is correct because transit gateway route tables provide the required segmentation and

routing control, allowing the company to decide which VPC attachments can communicate with each other and with the on-premises network. This design is the standard scalable approach for multi-account, multi-VPC connectivity with centralized control. See AWS documentation for [sharing a transit gateway by using AWS RAM](#) and [Transit Gateway route tables](#).

QUESTION NO: 130

A company is running an application in the AWS Cloud. The company's security team must approve the creation of all new IAM users. When a new IAM user is created, all access for the user must be removed automatically. The security team must then receive a notification to approve the user. The company has a multi-Region AWS CloudTrail trail in the AWS account.

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

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule. Define a pattern with the detail-type value set to AWS API Call via CloudTrail and an eventName of CreateUser.
- B. Configure CloudTrail to send a notification for the CreateUser event to an Amazon Simple Notification Service (Amazon SNS) topic.
- C. Invoke a container that runs in Amazon Elastic Container Service (Amazon ECS) with AWS Fargate technology to remove access
- D. Invoke an AWS Step Functions state machine to remove access.
- E. Use Amazon Simple Notification Service (Amazon SNS) to notify the security team.
- F. Use Amazon Pinpoint to notify the security team.

ANSWER: A D E

Explanation:

The correct combination is to create an Amazon EventBridge rule that matches CloudTrail management events for the IAM CreateUser API call, invoke an AWS Step Functions state machine to remove access, and use Amazon Simple Notification Service to notify the security team. CloudTrail records IAM API activity, and EventBridge can match those CloudTrail events in near real time by using an event pattern with detail-type set to AWS API Call via CloudTrail and eventName set to CreateUser. That provides the automated trigger whenever a new IAM user is created. Step Functions is appropriate for the remediation workflow because it can orchestrate AWS SDK service integrations and perform the required IAM cleanup actions, such as removing policies, group membership, access keys, or other permissions associated with the new user before approval. After remediation, Amazon Simple Notification Service provides a straightforward managed notification mechanism for emailing or otherwise alerting the security team so they can review and approve the new user. This approach uses event-driven automation and managed services, which aligns with AWS best practices for security response automation. See [Amazon EventBridge events delivered through CloudTrail](#) and [AWS Step Functions service integrations](#).

QUESTION NO: 131

A company's factory and automation applications are running in a single VPC. More than 20 applications run on a combination of Amazon EC2, Amazon Elastic Container Service (Amazon ECS), and Amazon RDS. The company has software engineers spread across three teams. One of the three teams owns each application, and each team is responsible for the cost and performance of all of its applications. Team resources have tags that represent their application and team. The teams use IAM access for daily activities. The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports. Which combination of actions will meet these requirements? (Choose three.)

- A. Activate the user-defined cost allocation tags that represent the application and the team.
- B. Activate the AWS generated cost allocation tags that represent the application and the team.
- C. Create a cost category for each application in Billing and Cost Management.

D. Activate IAM access to Billing and Cost Management.

E. Create a cost budget.

F. Enable Cost Explorer.

ANSWER: A C F

Explanation:

Activate the user-defined cost allocation tags that represent the application and the team, Create a cost category for each application in Billing and Cost Management, and Enable Cost Explorer are the correct combination. User-defined cost allocation tags are the right foundation because the company already tags resources by application and team; after activation, AWS can include those tag keys in billing and cost analysis data so charges can be attributed to the correct ownership dimensions. AWS Cost Categories can then be used to organize and group costs into meaningful business views, such as application-oriented groupings, which makes chargeback, showback, and reporting easier across many workloads and teams. Cost Explorer provides the reporting capability required here: it can analyze historical cost and usage, group and filter by dimensions such as tags and cost categories, compare prior periods, and generate forecasts based on historical usage patterns. Together, these features support monthly attribution, 12-month historical comparisons, and forward-looking cost forecasts. See AWS documentation for [cost allocation tags](#) and [Cost Explorer](#).

QUESTION NO: 132

A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing. Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

A. Enable Aurora Auto Scaling for Aurora Replicas. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

B. Enable Aurora Auto Scaling for Aurora writers. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.

C. Enable Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing and sticky sessions enabled.

D. Enable Aurora Scaling for Aurora writers. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

ANSWER: C

Explanation:

Enable Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing and sticky sessions enabled. is correct because it addresses both scaling requirements in the architecture. For the database tier, Aurora Auto Scaling can automatically adjust the number of Aurora Replicas based on workload, helping the Aurora PostgreSQL cluster scale read capacity as user demand grows. This is the appropriate scaling mechanism for Aurora PostgreSQL because read replicas can be added or removed automatically while the cluster maintains a single primary writer endpoint for writes. For the application tier, an Application Load Balancer integrates well with EC2 Auto Scaling and can distribute HTTP/HTTPS traffic across multiple application instances. Because the application is stateful, enabling sticky sessions helps maintain session affinity so that requests from the same user are consistently routed to the same target, improving user experience during horizontal scaling. The round robin routing algorithm is a valid Application Load Balancer routing approach for distributing requests across healthy targets. Together, Aurora Replica Auto Scaling and an Application Load Balancer with sticky sessions allow the web/application layer and database read tier to scale while preserving session consistency. See [Aurora Auto Scaling with Aurora Replicas](#) and [sticky sessions for Application Load Balancers](#).

QUESTION NO: 133

A solutions architect is planning to migrate critical Microsoft SQL Server databases to AWS. Because the databases are legacy systems, the solutions architect will move the databases to a modern data architecture. The solutions architect must migrate the databases with near-zero downtime. Which solution will meet these requirements?

- A.** Use AWS Application Migration Service and the AWS Schema Conversion Tool (AWS SCT). Perform an in-place upgrade before the migration. Export the migrated data to Amazon Aurora Serverless after cutover. Repoint the applications to Amazon Aurora.
- B.** Use AWS Database Migration Service (AWS DMS) to rehost the database. Set Amazon S3 as a target. Set up change data capture (CDC) replication. When the source and destination are fully synchronized, load the data from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB instance.
- C.** Use native database high availability tools. Connect the source system to an Amazon RDS for Microsoft SQL Server DB instance. Configure replication accordingly. When data replication is finished, transition the workload to an Amazon RDS for Microsoft SQL Server DB instance.
- D.** Use AWS Application Migration Service. Rehost the database server on Amazon EC2. When data replication is finished, detach the database and move the database to an Amazon RDS for Microsoft SQL Server DB instance. Reattach the database and then cut over all networking.
- E.** Use AWS SCT to convert the SQL Server schema to Amazon Aurora. Use AWS DMS with full load and CDC replication to migrate the data, then cut over the applications after the target database is synchronized.

ANSWER: E

Explanation:

Using AWS Schema Conversion Tool (AWS SCT) with AWS Database Migration Service (AWS DMS) is the correct approach for modernizing legacy Microsoft SQL Server databases while keeping downtime very low. AWS SCT helps assess and convert the source SQL Server schema, database code, and objects to a target engine such as Amazon Aurora PostgreSQL-Compatible Edition or Amazon Aurora MySQL-Compatible Edition. AWS DMS can then perform an initial full load of the existing data and continue replicating ongoing changes from the SQL Server source by using change data capture (CDC). This lets the target Aurora database stay synchronized while the source application remains online. After validation, the team can schedule a brief cutover window, stop writes to the source, allow final CDC changes to apply, and repoint the applications to Aurora. This pattern is commonly used for heterogeneous database migrations and supports the near-zero downtime requirement. See the AWS documentation for [AWS Database Migration Service](#) and [AWS Schema Conversion Tool](#).

QUESTION NO: 134

A company has an application that stores user-uploaded videos in an Amazon S3 bucket that uses S3 Standard storage. Users access the videos frequently in the first 180 days after the videos are uploaded. Access after 180 days is rare. Named users and anonymous users access the videos. Most of the videos are more than 100 MB in size. Users often have poor internet connectivity when they upload videos, resulting in failed uploads. The company uses multipart uploads for the videos. A solutions architect needs to optimize the S3 costs of the application. Which combination of actions will meet these requirements? (Choose two.)

- A.** Configure the S3 bucket to be a Requester Pays bucket.
- B.** Use S3 Transfer Acceleration to upload the videos to the S3 bucket.
- C.** Create an S3 Lifecycle configuration to expire incomplete multipart uploads 7 days after initiation.
- D.** Create an S3 Lifecycle configuration to transition objects to S3 Glacier Instant Retrieval after 1 day.
- E.** Create an S3 Lifecycle configuration to transition objects to S3 Standard-infrequent Access (S3 Standard- IA) after 180 days.

ANSWER: C E

Explanation:

A lifecycle rule that expires incomplete multipart uploads 7 days after initiation is correct because failed or abandoned multipart uploads can leave uploaded parts stored in Amazon S3, and those parts continue to incur storage charges until the upload is completed or aborted. Amazon S3 Lifecycle can automatically delete those incomplete multipart upload parts after a configured number of days, which directly addresses the failed uploads caused by poor connectivity. See the AWS documentation on [aborting incomplete multipart uploads with a lifecycle configuration](#).

A lifecycle rule that transitions objects to S3 Standard-Infrequent Access after 180 days is also correct because the videos are accessed frequently during the first 180 days and rarely afterward. S3 Standard-IA is designed for data that is accessed less frequently but still requires rapid access when needed. Since most videos are larger than 100 MB, they comfortably exceed the small-object threshold where infrequent access storage classes are typically cost effective. AWS documents S3 Lifecycle transitions and storage class considerations in the [S3 Lifecycle transition considerations](#).