PassQueen
a.k.a dumpsqueen.com

AWS Certified Solutions Architect - Professional

Amazon AWS SAP-C02

Version Demo

Total Demo Questions: 157

Total Premium Questions: 1571

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Design Solutions for Organizational Complexity	389
Topic 2, Design for New Solutions	467
Topic 3, Continuous Improvement for Existing Solutions	441
Topic 4, Accelerate Workload Migration and Modernization	274
Total	1571

QUESTION NO: 1

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 transit gateways in the applicable AWS Regions. Attach all the company's VPCs and relevant subnets to the transit gateways, and use transit gateway peering for inter-Region connectivity.

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:

Creating AWS Transit Gateways in the applicable Regions, attaching the VPCs, and connecting the regional transit gateways by using transit gateway peering provides scalable any-to-any private connectivity among the company's VPCs. Transit Gateway is a hub-and-spoke routing service that can be shared across AWS accounts through AWS Resource Access Manager. It avoids the rapidly growing number of connections and route-management tasks that would result from a full mesh of VPC peering connections. Transit gateway route tables can also be used to centrally control which attachment prefixes are reachable. Because a transit gateway is a Regional resource, VPCs in different Regions use their respective Regional transit gateways, with inter-Region transit gateway peering to carry traffic between them. See AWS Transit Gateway documentation.

Creating a Network Load Balancer for the license-validation workload and publishing it through an AWS PrivateLink endpoint service gives each customer VPC private, consumer-initiated access to that specific management service. Customers create interface endpoints to reach the endpoint service, while the provider retains control of the service and does not expose general network access to the management VPC. This is well suited to a SaaS provider serving an increasing number of isolated customer VPCs, including VPCs in other accounts, with minimal per-customer networking administration. See AWS PrivateLink documentation.

QUESTION NO: 2

A solutions architect needs to copy data from an Amazon S3 bucket m an AWS account to a new S3 bucket in a new AWS account. The solutions architect must implement a solution that uses the AWS CLI.

Which combination of steps will successfully copy the data? (Choose three.)

A. Create a bucket policy to allow the source bucket to list its contents and to put objects and set object ACLs in the destination bucket. Attach the bucket policy to the destination bucket.

B. Create a bucket policy to allow a user In the destination account to list the source bucket's contents and read the source bucket's objects. Attach the bucket policy to the source bucket.

C. Create an IAM policy in the source account. Configure the policy to allow a user In the source account to list contents and get objects In the source bucket, and to list contents, put objects, and set object ACLs in the destination bucket. Attach the policy to the user _

D. Create an IAM policy in the destination account. Configure the policy to allow a user In the destination account to list contents and get objects In the source bucket, and to list contents, put objects, and set objectACLs in the destination bucket. Attach the policy to the user.

E. Run the aws s3 sync command as a user in the source account. Specify' the source and destination buckets to copy the data.

F. Run the aws s3 sync command as a user in the destination account. Specify' the source and destination buckets to copy the data.

ANSWER: B D F

Explanation:

Cross-account S3 copying requires the IAM principal that runs the AWS CLI command to have identity-based permissions and to be authorized by the source bucket's resource-based policy. "Create a bucket policy to allow a user In the destination account to list the source bucket's contents and read the source bucket's objects. Attach the bucket policy to the source bucket." grants the destination-account principal access at the source resource boundary. The permissions to list the bucket and retrieve its objects allow the CLI operation to discover and read the source data.

"Create an IAM policy in the destination account. Configure the policy to allow a user In the destination account to list contents and get objects In the source bucket, and to list contents, put objects, and set objectACLs in the destination bucket. Attach the policy to the user." supplies the required identity permissions for that same principal, including writing the copied objects to the destination bucket. "Run the aws s3 sync command as a user in the destination account. Specify' the source and destination buckets to copy the data." then performs the transfer with a principal that can read from the source and write to the destination. This also results in objects being written by the destination-account principal. See Amazon S3 cross-account access guidance and the AWS CLI s3 sync reference.

QUESTION NO: 3

A company needs to store and process image data that will be uploaded from mobile devices using a custom mobile app. Usage peaks between 8 AM and 5 PM on weekdays, with thousands of uploads per minute. The app is rarely used at any other time. A user is notified when image processing is complete.

Which combination of actions should a solutions architect take to ensure image processing can scale to handle the load? (Select THRE

E.

)

A. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon MQ queue.

B. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue.

C. Invoke an AWS Lambda function to perform image processing when a message is available in the queue.

D. Invoke an S3 Batch Operations job to perform image processing when a message is available in the queue

E. Send a push notification to the mobile app by using Amazon Simple Notification Service (Amazon SNS) when processing is complete.

F. Send a push notification to the mobile app by using Amazon Simple Email Service (Amazon SES) when processing is complete.

ANSWER: B C E

Explanation:

"Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue," "Invoke an AWS Lambda function to perform image processing when a message is available in the queue," and "Send a push notification to the mobile app by using Amazon Simple Notification Service (Amazon SNS) when processing is complete" form a serverless, highly scalable asynchronous workflow. Direct uploads to Amazon S3 avoid routing high-volume image data through application servers and allow S3 to absorb the large, bursty upload rate. S3 event notifications can publish object-created events to an SQS standard queue, which

decouples ingestion from processing and durably buffers work during weekday peaks. Lambda can consume messages from SQS through an event source mapping and automatically scales concurrent processing in response to queue backlog, subject to configured concurrency and account quotas. After a successful processing operation, Amazon SNS can deliver a mobile push notification through supported mobile push platforms, allowing the app user to be informed without maintaining a persistent connection to the processing service. This architecture also aligns ongoing compute cost with actual demand because the workload is rarely active outside business hours. See the Amazon S3 event notification documentation and the AWS Lambda with Amazon SQS documentation.

QUESTION NO: 4

A company is developing an application that will display financial reports. The company needs a solution that can store financial Information that comes from multiple systems. The solution must provide the reports through a web interface and must serve the data will less man 500 milliseconds or latency to end users. The solution also must be highly available and must have an RTO or 30 seconds.

Which solution will meet these requirements?

A. Use an Amazon Redshift cluster to store the data. Use a static website hosted on Amazon S3 with backend APIs served by an Amazon Elastic Kubernetes Service (Amazon EKS) cluster to provide the reports to the application.

B. Use Amazon S3 to store the data. Use Amazon Athena to provide the reports to the application. Use AWS App Runner to serve the application to view the reports.

C. Use Amazon DynamoDB to store the data. Use an embedded Amazon QuickSight dashboard with direct query datasets to provide the reports to the application.

D. Use Amazon Keyspaces (for Apache Cassandra) to store the data. Use AWS Elastic Beanstalk to provide the reports to the application.

ANSWER: C

Explanation:

Use Amazon DynamoDB to store the data and an embedded Amazon QuickSight dashboard with direct query datasets to provide the reports to the application. DynamoDB is designed for consistently low-latency operational data access, with single-digit-millisecond performance at scale. This gives the reporting application substantial margin to meet the end-to-end latency target of less than 500 milliseconds, assuming the data model and access patterns are designed appropriately. DynamoDB is also a managed, multi-AZ service: data is synchronously replicated across multiple Availability Zones in an AWS Region, so an Availability Zone failure does not require the company to perform database recovery or infrastructure failover. That managed resilience supports the 30-second recovery objective for the application's data layer. QuickSight supplies the browser-based reporting experience, and embedding its dashboards lets the company place interactive financial reports directly into the application's web interface without developing visualization capabilities from scratch. Direct query datasets retrieve current source data when the dashboard is queried, which is appropriate when reports must reflect updates arriving from multiple systems. Review the Amazon DynamoDB Developer Guide for its performance and high-availability design, and the Amazon QuickSight embedded analytics documentation for dashboard embedding.

QUESTION NO: 5

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 1AM 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 A.

B. In Account B.

C. In Account A, add an S3 bucket policy that grants User_DataProcessor in Account B permission to access the required bucket objects.

D. In Account B, attach an IAM policy to User_DataProcessor that allows the required Amazon S3 actions on the bucket and its objects in Account A.

ANSWER: C D

Explanation:

Cross-account access to an Amazon S3 bucket requires permission from both sides of the access request. The bucket owner in Account A must use an S3 bucket policy to name the external IAM principal, `arn:aws:iam::Account-B-ID:user/User_DataProcessor`, as a principal and grant only the required S3 actions, such as `s3:GetObject`, on the required bucket and object ARNs. This resource-based policy establishes that the bucket permits the user from the partner account to access the files.

The business partner must also attach an identity-based IAM policy to User_DataProcessor in Account B. That policy authorizes the user to make the required S3 API calls against the Account A bucket and its objects. AWS evaluates the user's identity-based permissions together with the bucket policy; the request succeeds when both policies allow the requested action and no applicable explicit deny exists. For read access, the policy commonly includes `s3:ListBucket` on the bucket ARN when listing is required and `s3:GetObject` on the object ARN or permitted prefix.

See the AWS guidance for cross-account Amazon S3 access and the Amazon S3 bucket policy examples.

QUESTION NO: 6

A company manages hundreds of AWS accounts centrally in an organization In AWS Organizations. The company recently started to allow product teams to create and manage their own S3 access points in their accounts. The S3 access points can be accessed only within VPCs. not on the internet.

What is the MOST operationally efficient way to enforce this requirement?

A. Set the S3 access point resource policy to deny the `s3:CreateAccessPoint` action unless the `s3:AccessPointNetworkOrigin` condition key evaluates to `VPC`.

B. Create an SCP at the root level in the organization to deny the `s3:CreateAccessPoint` action unless the `s3:AccessPointNetworkOrigin` condition key evaluates to `VPC`.

C. Use AWS CloudFormation StackSets to create a new IAM policy in each AWS account that allows the `s3:CreateAccessPoint` action only if the `s3:AccessPointNetworkOrigin` condition key evaluates to `VPC`.

D. Set the S3 bucket policy to deny the `s3:CreateAccessPoint` action unless the `s3:AccessPointNetworkOrigin` condition key evaluates to `VPC`.

ANSWER: B

Explanation:

Create an SCP at the root level in the organization to deny the `s3:CreateAccessPoint` action unless the `s3:AccessPointNetworkOrigin` condition key evaluates to `VPC`. This creates a centralized preventive guardrail for every account governed by the organization, while allowing product teams to retain the ability to create and administer compliant access points in their own accounts. Amazon S3 access points support a network origin of either `Internet` or `VPC`. The `s3:AccessPointNetworkOrigin` condition key can be evaluated during access point creation to require the VPC origin. An explicit SCP deny takes precedence over IAM permissions that might otherwise permit access-point creation, so account-level administrators cannot inadvertently or intentionally create an internet-origin access point. Applying the SCP at the organization root delivers the policy once and consistently protects the large account estate, avoiding repeated deployment and ongoing maintenance of separate identity policies in each account. SCPs define the maximum available permissions for principals in member accounts and are designed for precisely this kind of organization-wide permission boundary. See the Amazon S3 access point policy documentation and the AWS Organizations SCP documentation.

QUESTION NO: 7

A medical company is running an application in the AWS Cloud. The application simulates the effect of medical drugs in development.

The application consists of two parts configuration and simulation The configuration part runs in AWS Fargate containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The simulation part runs on large, compute optimized Amazon EC2 instances Simulations can restart if they are interrupted

The configuration part runs 24 hours a day with a steady load. The simulation part runs only for a few hours each night with a variable load. The company stores simulation results in Amazon S3, and researchers use the results for 30 days. The company must store simulations for 10 years and must be able to retrieve the simulations within 5 hours

Which solution meets these requirements MOST cost-effectively?

A. Purchase an EC2 Instance Savings Plan to cover the usage for the configuration part Run the simulation part by using EC2 Spot Instances Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Intelligent-Tiering

B. Purchase an EC2 Instance Savings Plan to cover the usage for the configuration part and the simulation part Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier

C. Purchase Compute Savings Plans to cover the usage for the configuration part Run the simulation part by using EC2 Spot instances Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier

D. Purchase Compute Savings Plans to cover the usage for the configuration part Purchase EC2 Reserved Instances for the simulation part Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier Deep Archive

ANSWER: C

Explanation:

Purchase Compute Savings Plans to cover the usage for the configuration part Run the simulation part by using EC2 Spot instances Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier is the most cost-effective solution because it matches each component's cost model to its actual usage and availability requirements. The configuration workload runs continuously at a predictable level on AWS Fargate. Compute Savings Plans provide discounted pricing for committed, consistent compute use and apply to Fargate as well as eligible EC2 and Lambda usage. This makes them appropriate for the always-on ECS service without tying the commitment to a particular EC2 instance family or Region.

The nightly simulation workload is variable, runs for limited periods, and can restart following an interruption. EC2 Spot Instances are therefore suitable because they offer substantial savings for interruption-tolerant, flexible compute workloads. Simulation output remains actively used for its first 30 days, after which lifecycle management can archive it to S3 Glacier Flexible Retrieval, historically referred to as S3 Glacier. This archive class is designed for long-term, low-cost retention and provides standard retrievals typically within 3–5 hours, meeting the five-hour retrieval objective while supporting the required 10-year retention period. See AWS Savings Plans documentation and Amazon S3 storage classes documentation.

QUESTION NO: 8

A company uses Amazon S3 to store files and images in a variety of storage classes. The company's S3 costs have increased substantially during the past year.

A solutions architect needs to review data trends for the past 12 months and identity the appropriate storage class for the objects.

Which solution will meet these requirements?

A. Download AWS Cost and Usage Reports for the last 12 months of S3 usage. Review AWS Trusted Advisor recommendations for cost savings.

B. Use S3 storage class analysis. Import data trends into an Amazon QuickSight dashboard to analyze storage trends.

C. Use Amazon S3 Storage Lens. Upgrade the default dashboard to include advanced metrics for storage trends.

D. Use Access Analyzer for S3. Download the Access Analyzer for S3 report for the last 12 months. Import the csvfile to an Amazon QuickSight dashboard.

ANSWER: C

Explanation:

Use Amazon S3 Storage Lens. Upgrade the default dashboard to include advanced metrics for storage trends. This meets the requirement because S3 Storage Lens provides account- and organization-level visibility into S3 storage usage, activity, cost-efficiency, and data-protection trends. Its metrics can be viewed by dimensions such as bucket, prefix, AWS Region, and storage class, enabling the architect to identify where capacity is growing and which storage classes are contributing to the increased cost.

The Storage Lens default metrics retain historical data for only 14 days, which is insufficient for a 12-month review. Enabling the advanced metrics and recommendations tier provides up to 15 months of metric history. This allows analysis of the full required period and supports storage-class optimization decisions based on object age, access patterns, incomplete multipart uploads, noncurrent version counts, and other cost-efficiency indicators. The resulting trends help identify data that is appropriate for lifecycle transitions to lower-cost classes or for archival storage. S3 Storage Lens is specifically intended to provide these aggregated S3 analytics and optimization recommendations without requiring custom reporting infrastructure. See the Amazon S3 Storage Lens documentation and AWS guidance on viewing Storage Lens metrics.

QUESTION NO: 9

A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability.

Which service will meet the requirements for storing the session information in the MOST cost-effective way?

A. Amazon ElastiCache with the Memcached engine

B. Amazon S3

C. Amazon RDS MySQL

D. Amazon ElastiCache with the Redis engine

ANSWER: D

Explanation:

Amazon ElastiCache with the Redis engine is the appropriate session store because it keeps frequently accessed session state in memory, providing the very low read and write latency needed for a web application. The application servers can retrieve session data by key rather than relying on local instance storage. This makes the web tier stateless: instances can be added, replaced, or removed by an Auto Scaling group without losing a user's session or requiring session affinity to a particular server.

Redis supports expiration through time-to-live settings, which maps naturally to web-session lifecycles and automatically removes stale session records. It also provides replication groups, Multi-AZ capability, and automatic failover, allowing the session layer to remain available when an underlying node or Availability Zone has a failure. Because session data is normally transient and accessed at high frequency, a managed in-memory Redis deployment avoids the operational burden of managing software, replication, and failover while providing capacity that can be sized for the active session workload. AWS identifies ElastiCache as a solution for session management and documents Redis replication and automatic failover capabilities for highly available workloads. See AWS session management with caching and ElastiCache for Redis automatic failover.

QUESTION NO: 10

A company is running a web-crawling process on a list of target URLs to obtain training documents for machine learning training algorithms. A fleet of Amazon EC2 t2.micro instances pulls the target URLs from an Amazon Simple Queue Service (Amazon SQS) queue. The instances then write the result of the crawling algorithm as a .csv file to an Amazon Elastic File System (Amazon EFS) volume. The EFS volume is mounted on all instances of the fleet. A separate system adds the URLs to the SQS queue at infrequent rates. The instances crawl each URL in 10 seconds or less. Metrics indicate that some

instances are idle when no URLs are in the SQS queue. A solutions architect needs to redesign the architecture to optimize costs. Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

A. Use m5.8xlarge instances instead of t2.micro instances for the web-crawling process. Reduce the number of instances in the fleet by 50%.

B. Convert the web-crawling process into an AWS Lambda function. Configure the Lambda function to pull URLs from the SQS queue.

C. Modify the web-crawling process to store results in Amazon Neptune.

D. Modify the web-crawling process to store results in an Amazon Aurora Serverless MySQL instance.

E. Modify the web-crawling process to store results in Amazon S3.

ANSWER: B E

Explanation:

Convert the web-crawling process into an AWS Lambda function. Configure the Lambda function to pull URLs from the SQS queue. This is appropriate because the work is infrequent, event-driven, and completes within 10 seconds. Lambda integrates natively with Amazon SQS through an event source mapping. Lambda polls the queue and invokes the function only when messages are available, eliminating the cost of continuously running EC2 instances during idle periods. The function can process messages in batches and scales its concurrent processing in response to queue backlog, subject to configured concurrency controls. The crawler should also implement idempotent processing because Amazon SQS and Lambda event source mappings provide at-least-once delivery semantics.

Modify the web-crawling process to store results in Amazon S3. CSV crawl outputs are objects and can be written directly from each Lambda invocation to an S3 bucket. Amazon S3 provides highly durable, elastic object storage with pay-for-use pricing, avoiding the need to maintain a shared mounted file system for intermittent output. S3 is also well suited to retaining and organizing machine learning training documents through bucket prefixes, object metadata, encryption, lifecycle policies, and integration with downstream analytics and ML services. Together, SQS, Lambda, and S3 create a serverless design that removes idle compute costs and minimizes always-on infrastructure. See Using Lambda with Amazon SQS and Amazon S3 User Guide.

QUESTION NO: 11

A company hosts its primary API on AWS using Amazon API Gateway and AWS Lambda functions. Internal applications and external customers use this API. Some customers also use a legacy API hosted on a standalone EC2 instance.

The company wants to increase security across all APIs to prevent denial of service (DoS) attacks, check for vulnerabilities, and guard against common exploits.

What should a solutions architect do to meet these requirements?

A. Use AWS WAF to protect both APIs. Configure Amazon Inspector to analyze the legacy API. Configure Amazon GuardDuty to monitor for malicious attempts to access the APIs.

B. Use AWS WAF to protect the API Gateway API. Configure Amazon Inspector to analyze both APIs. Configure Amazon GuardDuty to block malicious attempts.

C. Use AWS WAF to protect the API Gateway API. Configure Amazon Inspector to analyze the legacy API. Configure Amazon GuardDuty to monitor for malicious attempts to access the APIs.

D. Use AWS WAF to protect the API Gateway API. Configure Amazon Inspector to protect the legacy API. Configure Amazon GuardDuty to block malicious attempts.

E. Associate AWS WAF with the API Gateway API and an Application Load Balancer in front of the EC2 API. Enable AWS Shield Advanced, use Amazon Inspector to scan the EC2 instance and Lambda functions, and use Amazon GuardDuty for threat monitoring.

ANSWER: E

Explanation:

Use AWS WAF with the API Gateway API stage and with an Application Load Balancer placed in front of the legacy EC2-hosted API. This provides web application-layer protection for both API entry points, including managed rule groups for common exploit patterns and rate-based rules that can limit excessive request rates. Placing the EC2 workload behind an Application Load Balancer also provides a supported AWS WAF integration point rather than exposing the instance directly to the internet.

Enable AWS Shield Advanced for the API Gateway API and the Application Load Balancer to add enhanced DDoS detection, mitigation, and response capabilities for these protected resources. Amazon Inspector should scan the EC2 instance and Lambda functions for software vulnerabilities and unintended network exposure. Amazon GuardDuty complements these controls by continuously monitoring AWS data sources for potentially malicious activity and producing security findings for investigation. Together, these services provide preventive edge protection, DDoS mitigation, vulnerability management, and ongoing threat detection across the API architecture. See the AWS WAF Developer Guide and the Amazon Inspector resource scanning documentation.

QUESTION NO: 12

A company receives application event data from multiple sources. The company uses Amazon Kinesis Data Firehose to deliver the events to an Amazon S3 data lake. Data analysts need to query the events by using standard SQL. New event attributes can be added over time, and the company wants to avoid managing database servers or manually updating table definitions whenever the event schema changes.

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

A. Configure an AWS Glue crawler to catalog the event data in Amazon S3 and update the AWS Glue Data Catalog.

B. Use Amazon Athena to query the cataloged event data in Amazon S3.

C. Create an Amazon RDS for PostgreSQL DB instance. Import every Firehose delivery file into relational tables.

D. Deploy an Amazon EMR cluster that runs continuously to create and maintain Hive table definitions.

E. Use Amazon OpenSearch Service as the only destination for the Firehose delivery stream.

ANSWER: A B

Explanation:

Configure an AWS Glue crawler to catalog the event data in Amazon S3 and update the AWS Glue Data Catalog. A crawler can inspect the data in S3, infer schemas, create tables, and update table metadata when new compatible attributes are discovered. This avoids manually maintaining table definitions as the event data evolves.

Use Amazon Athena to query the cataloged data in Amazon S3. Athena is a serverless query service that uses standard SQL and integrates with the AWS Glue Data Catalog. Analysts can query the Firehose-delivered data directly in S3 without provisioning, patching, or scaling database servers. The company can further reduce query cost by configuring Firehose to deliver partitioned and compressed data formats where appropriate. See AWS Glue crawlers and What is Amazon Athena?.

QUESTION NO: 13

A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.

Which set of actions should a solutions architect take to meet these requirements?

A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instances. Use Systems Manager to generate patch compliance reports.

B. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances. Use Amazon QuickSight integration with OpsWorks to generate patch compliance reports.

C. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to apply patches by scheduling an AWS Systems Manager patch remediation job. Use Amazon Inspector to generate patch compliance reports.

D. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

ANSWER: A

Explanation:

Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instances. Use Systems Manager to generate patch compliance reports. This provides a centralized patch-management and compliance solution for both Amazon EC2 instances and on-premises servers. EC2 instances can be managed as Systems Manager managed nodes when they have the SSM Agent, appropriate IAM permissions, and network connectivity to Systems Manager. On-premises servers can join the same management plane as hybrid and multicloud managed nodes by using a Systems Manager hybrid activation and installing the SSM Agent.

After the fleet is managed by Systems Manager, Patch Manager can assess patch status and apply patches according to patch baselines. The organization can schedule patching with maintenance windows or run patch operations on demand. Systems Manager records patch compliance information for managed nodes, including whether patches are compliant, missing, installed, or failed. This produces the required unified visibility across the mixed on-premises and AWS environment, rather than maintaining separate reporting mechanisms for each existing patching tool. Compliance data is available through Systems Manager Compliance and can be used for centralized operational reporting. See the AWS documentation for AWS Systems Manager Patch Manager and hybrid and multicloud managed nodes.

QUESTION NO: 14

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:

Installing and configuring the Amazon CloudWatch agent to send process metrics for the application provides health monitoring for the custom network analysis software itself, rather than monitoring only the underlying EC2 instance. The CloudWatch agent's `procstat` capability can publish metrics based on whether a named process is running, enabling application-level failure detection. AWS documents this process monitoring capability at CloudWatch agent procstat process metrics.

Creating a CloudWatch alarm for the custom process metric and publishing its state change to an Amazon SNS topic creates an automated, event-driven remediation path when the analysis software fails. An AWS Lambda function can

consume that notification, remove or mark the affected Auto Scaling instance unhealthy so that replacement occurs, and coordinate the route update after the replacement target is available. VPC route-table targets that identify a specific instance or network interface do not automatically change when an Auto Scaling group launches a replacement instance. The automation must update the route target explicitly, such as by calling the Amazon EC2 `ReplaceRoute` API, documented at ReplaceRoute API reference. Together, these steps detect the actual software failure and ensure traffic is redirected to the newly launched inspection instance.

QUESTION NO: 15

A company plans to migrate a legacy on-premises application to AWS. The application is a Java web application that runs on Apache Tomcat with a PostgreSQL database. The company does not have access to the source code but can deploy the application Java Archive (JAR) files. The application has increased traffic at the end of each month. Which solution will meet these requirements with the LEAST operational overhead?

A. Launch Amazon EC2 instances in multiple Availability Zones. Deploy Tomcat and PostgreSQL to all the instances by using Amazon Elastic File System (Amazon EFS) mount points. Use AWS Step Functions to deploy additional EC2 instances to scale for increased traffic.

B. Provision Amazon Elastic Kubernetes Service (Amazon EKS) in an Auto Scaling group across multiple AWS Regions. Deploy Tomcat and PostgreSQL in the container images. Use a Network Load Balancer to scale for increased traffic.

C. Refactor the Java application into Python-based containers. Use AWS Lambda functions for the application logic. Store application data in Amazon DynamoDB global tables. Use AWS Storage Gateway and Lambda concurrency to scale for increased traffic.

D. Use AWS Elastic Beanstalk to deploy the Tomcat servers with auto scaling in multiple Availability Zones. Store application data in an Amazon RDS for PostgreSQL database. Deploy Amazon CloudFront and an Application Load Balancer to scale for increased traffic.

ANSWER: D

Explanation:

Use AWS Elastic Beanstalk to deploy the Tomcat servers with auto scaling in multiple Availability Zones. Store application data in an Amazon RDS for PostgreSQL database. Deploy Amazon CloudFront and an Application Load Balancer to scale for increased traffic. This approach preserves the application's existing Java, Tomcat, and PostgreSQL architecture, avoiding a source-code-dependent redesign while minimizing infrastructure administration. Elastic Beanstalk is a managed application platform that provisions and operates the underlying compute capacity, load balancing, health monitoring, deployments, and Auto Scaling for a Java environment. The environment can span multiple Availability Zones, and its Auto Scaling configuration can add instances during the predictable end-of-month demand increase. An Application Load Balancer distributes application requests across the healthy instances, while CloudFront can cache eligible content at edge locations and reduce load on the application fleet. Amazon RDS for PostgreSQL provides a managed PostgreSQL-compatible database service, reducing routine operational tasks through automated backups, software maintenance capabilities, monitoring, and Multi-AZ deployment options. Together, these managed services allow the company to deploy the packaged application and retain the relational database model without operating Kubernetes, self-managed database servers, or custom scaling workflows. See Deploying Java applications with AWS Elastic Beanstalk and Amazon RDS for PostgreSQL documentation.

QUESTION NO: 16

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 (or 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:

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. This solution preserves the customers' existing connection target because an internet-facing, VPC-hosted AWS Transfer Family endpoint supports Elastic IP addresses. Customers can therefore continue to connect to the same public IP address with SFTP over SSH, avoiding DNS, client configuration, and workflow changes. The endpoint's security groups provide the required IP-based network access control.

AWS Transfer Family is a fully managed file-transfer service, so AWS operates the underlying highly available SFTP infrastructure rather than requiring the company to administer, patch, scale, or recover an SFTP application running on EC2. Amazon S3 provides durable, managed backend storage for transferred files and removes dependency on instance-local storage. Migrating the existing files to S3 completes the cutover while maintaining the established external access pattern. AWS documents the Elastic IP and security-group configuration for internet-facing VPC-hosted Transfer Family servers in Create a server in a VPC, and describes Transfer Family's managed SFTP service and S3 storage support in the AWS Transfer Family User Guide.

QUESTION NO: 17

A company is using multiple AWS accounts and has multiple DevOps teams running production and non-production workloads in these accounts. The company would like to centrally-restrict access to some of the AWS services that the DevOps teams do not use. The company decided to use AWS Organizations and successfully invited all AWS accounts into the Organization. They would like to allow access to services that are currently in-use and deny a few specific services. Also they would like to administer multiple accounts together as a single unit. What combination of steps should the solutions architect take to satisfy these requirements? (Choose three.)

A. Use a Deny list strategy.

B. Review the Access Advisor in AWS IAM to determine services recently used

C. Review the AWS Trusted Advisor report to determine services recently used.

D. Remove the default FullAWSAccess SCP.

E. Define organizational units (OUs) and place the member accounts in the OUs.

F. Remove the default DenyAWSAccess SCP.

ANSWER: A B E

Explanation:

Using **Use a Deny list strategy.** fits the stated governance model because the company wants accounts to retain access to broadly available AWS services while centrally blocking only a limited set of unwanted services. With this strategy, the organization retains the default `FullAWSAccess` service control policy (SCP) and attaches additional SCPs containing explicit deny statements for the services or actions that must be prohibited. SCPs establish the maximum permissions that IAM principals in affected member accounts can receive, enabling centralized preventative controls.

Review the Access Advisor in AWS IAM to determine services recently used provides service-last-accessed data for IAM entities and policies. This helps the company validate its understanding of current service usage before enforcing restrictions, reducing the risk of disrupting production or non-production workloads that rely on a service.

Define organizational units (OUs) and place the member accounts in the OUs. enables accounts to be grouped into manageable administrative units, such as production and non-production OUs. The organization can then attach SCPs to an OU and have those controls inherited by accounts within that unit, while applying different controls to other OUs as needed. See the AWS documentation for SCP strategies and IAM Access Advisor.

QUESTION NO: 18

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: A C

Explanation:

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. This provides separate, application-specific access-control entry points to the shared data lake. Access point policies can grant only the required bucket and object actions and prefixes for each workload, supporting least privilege. An access point with a VPC network origin accepts requests only from the specified VPC, and the bucket policy can require requests to use an access point rather than the bucket's direct endpoint.

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. An S3 gateway endpoint provides private routing from resources in the VPC to Amazon S3 without requiring an internet gateway, NAT gateway, or public IP address. Associating the endpoint with the relevant route tables ensures application traffic uses that private path. The endpoint policy adds another policy layer that can limit use of the endpoint to the intended S3 access point. See Amazon S3 access points and Gateway endpoints for Amazon S3.

QUESTION NO: 19

A company has an Amazon VPC that is divided into a public subnet and a pnvate 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 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

Onty 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 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 required rules are **An inbound rule for port 80 from source 10.0 0 0/24** and **An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24**. The Application Load Balancer nodes are located in the public subnet, whose CIDR range is 10.0.0.0/24. When the load balancer forwards an HTTP request to a web server target, the request enters the private subnet on the target listener port, TCP port 80. Restricting that inbound NACL rule to the public subnet CIDR permits only the required load-balancer-to-target request path between these subnets.

Network ACLs are stateless. Consequently, the private subnet NACL must separately allow the return path for each permitted TCP flow. The load balancer initiates its connection to a target from an ephemeral source port. Responses sent by the EC2 web server therefore have a destination port in the ephemeral range. Allowing outbound TCP ports 1024 through 65535 to 10.0.0.0/24 enables response packets to return to the load balancer nodes while maintaining the requested subnet-level restriction. AWS recommends accounting for ephemeral ports in both directions whenever using custom network ACLs. See the AWS documentation on network ACLs and custom network ACL rules and ephemeral ports.

QUESTION NO: 20

A company is using AWS Control Tower to manage AWS accounts in an organization in AWS Organizations. The company has an OU that contains accounts. The company must prevent any new or existing Amazon EC2 instances in the OU's accounts from gaining a public IP address. Which solution will meet these requirements?

A. Configure all instances in each account in the OU to use AWS Systems Manager. Use a Systems Manager Automation runbook to prevent public IP addresses from being attached to the instances.

B. Implement the AWS Control Tower proactive control to check whether instances in the OU's accounts have a public IP address. Set the AssociatePublicIpAddress property to False. Attach the proactive control to the OU.

C. Create an SCP that prevents the launch of instances that have a public IP address. Additionally, configure the SCP to prevent the attachment of a public IP address to existing instances. Attach the SCP to the OU.

D. Create an AWS Config custom rule that detects instances that have a public IP address. Configure a remediation action that uses an AWS Lambda function to detach the public IP addresses from the instances.

ANSWER: C

Explanation:

Create an SCP that prevents the launch of instances that have a public IP address. Additionally, configure the SCP to prevent the attachment of a public IP address to existing instances. Attach the SCP to the OU. This is the appropriate preventive, organization-wide solution. A service control policy attached to an organizational unit establishes permission boundaries for every member account in that OU, including accounts governed through AWS Control Tower. SCP explicit denies apply even if an IAM policy in a member account would otherwise allow the operation.

The policy can deny `ec2:RunInstances` when the request specifies public IPv4 assignment, using the `ec2:AssociatePublicIpAddress` condition key. This prevents launches that explicitly request a public address and launches that would receive one through the relevant request setting. It can also deny `ec2:AssociateAddress`, preventing an Elastic IP address from being associated with an existing instance or its network interface. Together, these controls prevent the relevant public-IP assignment paths for newly launched and already-running instances, before exposure occurs. SCPs are therefore suited to the requirement to prevent, rather than merely detect and remediate, noncompliant configurations. AWS documents SCP evaluation and their organization-level permission effect in the AWS Organizations SCP documentation. The supported EC2 actions and request condition keys are listed in the Amazon EC2 service authorization reference.

QUESTION NO: 21

A utility company collects usage data from smart meters every 5 minutes. Data is sent to API Gateway, processed by Lambda, and stored in DynamoDB. As usage increased, Lambda durations increased and DynamoDB PUTs failed with ProvisionedThroughputExceededException. Lambda also experiences TooManyRequestsException errors.

Which combination of changes will resolve these issues? (Select TWO.)

A. Increase the write capacity units to the DynamoDB table.

B. Increase the memory available to the Lambda functions.

C. Increase the payload size from the smart meters.

D. Stream the data into an Amazon Kinesis data stream from API Gateway and process the data in batches.

E. Collect data in an Amazon SQS FIFO queue, which triggers a Lambda function to process each message.

ANSWER: A D

Explanation:

Increasing the write capacity units to the DynamoDB table directly addresses `ProvisionedThroughputExceededException` errors. In provisioned capacity mode, DynamoDB requires sufficient write capacity units to sustain the table's incoming write rate. Capacity should be sized for the number and size of meter records written per second, with consideration for indexes because writes to a global secondary index also consume write capacity. This allows the table to accept the increased ingestion workload without throttling.

Streaming the data into an Amazon Kinesis data stream from API Gateway and processing the data in batches decouples request ingestion from database processing. Kinesis durably absorbs bursty meter submissions, while Lambda reads records in batches rather than invoking once for every individual meter event. Batch processing reduces invocation overhead and allows Lambda concurrency to be managed more efficiently, helping prevent Lambda request throttling represented by `TooManyRequestsException`. The buffered, scalable ingestion path also enables the processing rate to be aligned with the DynamoDB write capacity, rather than requiring every API request to synchronously complete a DynamoDB write.

See the DynamoDB provisioned throughput documentation and AWS Lambda integration with Kinesis.

QUESTION NO: 22

A solutions architect is creating an application that stores objects in an Amazon S3 bucket The solutions architect must deploy the application in two AWS Regions that will be used simultaneously The objects in the two S3 buckets must remain synchronized with each other.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Select THREE)

A. Create an S3 Multi-Region Access Point. Change the application to refer to the Multi-Region Access Point

B. Configure two-way S3 Cross-Region Replication (CRR) between the two S3 buckets

C. Modify the application to store objects in each S3 bucket.

D. Create an S3 Lifecycle rule for each S3 bucket to copy objects from one S3 bucket to the other S3 bucket.

E. Enable S3 Versioning for each S3 bucket

F. Configure an event notification for each S3 bucket to invoke an AVVS Lambda function to copy objects from one S3 bucket to the other S3 bucket.

ANSWER: A B E

Explanation:

Create an S3 Multi-Region Access Point. Change the application to refer to the Multi-Region Access Point provides one global S3 endpoint for the application rather than requiring application code to select and manage individual regional bucket endpoints. S3 Multi-Region Access Points are designed for multi-Region S3 architectures and can route requests to the closest active bucket, reducing application and operational complexity.

Configure two-way S3 Cross-Region Replication (CRR) between the two S3 buckets establishes managed replication rules in both directions. Therefore, objects uploaded in either Region are replicated to the bucket in the other Region, supporting simultaneous use of both Regions without custom object-copy workflows. Replication is asynchronous, so the design should account for replication status where an immediate cross-Region read is required.

Enable S3 Versioning for each S3 bucket is a prerequisite for S3 Replication. Versioning allows S3 to identify and replicate object versions consistently between the replication source and destination buckets. These managed S3 capabilities together provide the required active-active multi-Region storage pattern with the least ongoing operational effort. See the Amazon S3 Multi-Region Access Points documentation and the Amazon S3 Replication documentation.

QUESTION NO: 23

A solutions architect is creating an application that stores objects in an Amazon S3 bucket. The solutions architect must deploy the application in two AWS Regions that will be used simultaneously. The objects in the two S3 buckets must remain synchronized with each other. Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose three.)

A. Create an S3 Multi-Region Access Point Change the application to refer to the Multi-Region Access Point

B. Configure two-way S3 Cross-Region Replication (CRR) between the two S3 buckets

C. Modify the application to store objects in each S3 bucket

D. Create an S3 Lifecycle rule for each S3 bucket to copy objects from one S3 bucket to the other S3 bucket

E. Enable S3 Versioning for each S3 bucket

F. Configure an event notification for each S3 bucket to invoke an AWS Lambda function to copy objects from one S3 bucket to the other S3 bucket

ANSWER: A B E

Explanation:

Create an S3 Multi-Region Access Point Change the application to refer to the Multi-Region Access Point, **Configure two-way S3 Cross-Region Replication (CRR) between the two S3 buckets**, and **Enable S3 Versioning for each S3 bucket** provide a managed active-active multi-Region S3 design. A Multi-Region Access Point exposes a single global endpoint for the application while supporting requests to buckets in multiple Regions. This removes the need for the application to select and manage regional S3 endpoints itself. Two-way S3 Cross-Region Replication establishes a replication rule in each direction, so writes made in either regional bucket can be asynchronously replicated to the other bucket. This managed replication approach avoids maintaining custom copy services, queues, event handlers, or reconciliation logic. S3 Versioning must be enabled on both source and destination buckets because it is a prerequisite for Amazon S3 Replication. Versioning gives S3 the object-version semantics necessary to track replicated updates and delete markers correctly. Together, these features use native S3 capabilities to provide a common access endpoint and synchronization across both Regions with minimal operational management. See the AWS documentation for Multi-Region Access Points and Replicating objects within and across Regions.

QUESTION NO: 24

A company that provides image storage services wants to deploy a customer-facing solution to AWS. Millions of individual customers will use the solution. The solution will receive batches of large image files, resize the files, and store the files in an Amazon S3 bucket for up to 6 months. The solution must handle significant variance in demand. The solution must also be reliable at enterprise scale and have the ability to rerun processing jobs in the event of failure. Which solution will meet these requirements MOST cost-effectively?

A. Use AWS Step Functions to process the S3 event that occurs when a user stores an image. Run an AWS Lambda function that resizes the image in place and replaces the original file in the S3 bucket. Create an S3 Lifecycle expiration policy to expire all stored images after 6 months.

B. Use Amazon EventBridge to process the S3 event that occurs when a user uploads an image. Run an AWS Lambda function that resizes the image in place and replaces the original file in the S3 bucket. Create an S3 Lifecycle expiration policy to expire all stored images after 6 months.

C. Use S3 Event Notifications to invoke an AWS Lambda function when a user stores an image. Use the Lambda function to resize the image in place and to store the original file in the S3 bucket. Create an S3 Lifecycle policy to move all stored images to S3 Standard- Infrequent Access (S3 Standard-IA) after 6 months.

D. Use Amazon Simple Queue Service (Amazon SQS) to process the S3 event that occurs when a user stores an image. Run an AWS Lambda function that resizes the image and stores the resized file in an S3 bucket that uses S3 Standard-Infrequent Access (S3 Standard-IA). Create an S3 Lifecycle expiration policy to delete all stored images after 6 months.

ANSWER: D

Explanation:

Use Amazon Simple Queue Service (Amazon SQS) to process the S3 event that occurs when a user stores an image. Run an AWS Lambda function that resizes the image and stores the resized file in an S3 bucket that uses S3 Standard-Infrequent Access (S3 Standard-IA). Create an S3 Lifecycle expiration policy to delete all stored images after 6 months. This architecture decouples image uploads from processing, which is essential when customer demand arrives in unpredictable, high-volume batches. Amazon S3 event notifications can send object-created events to an SQS queue, where messages durably buffer work until Lambda consumers process them. Lambda scales the processing fleet automatically without requiring the company to maintain continuously running servers, making it a cost-effective fit for variable demand. SQS also supports reliable asynchronous processing: failed messages can become visible for retry, and a dead-letter queue can retain messages that repeatedly fail so they can be inspected and reprocessed. Storing the output in S3 Standard-IA reduces cost for data expected to be retained for months with infrequent access, subject to its applicable storage-duration and retrieval pricing. Finally, an expiration lifecycle rule enforces the stated maximum six-month retention period. See the Amazon S3 event notification documentation and the AWS Lambda documentation for Amazon SQS event source mappings.

QUESTION NO: 25

A company has a photo sharing social networking application. To provide a consistent experience for users, the company performs some image processing on the photos uploaded by users before publishing on the application. The image processing is implemented using a set of Python libraries.

The current architecture is as follows:

• The image processing Python code runs in a single Amazon EC2 instance and stores the processed images in an Amazon S3 bucket named ImageBucket.

• The front-end application, hosted in another bucket, loads the images from ImageBucket to display to users.

With plans for global expansion, the company wants to implement changes in its existing architecture to be able to scale for increased demand on the application and reduce management complexity as the application scales.

Which combination of changes should a solutions architect make? (Select TWO.)

A. Place the image processing EC2 instance into an Auto Scaling group.

B. Use AWS Lambda to run the image processing tasks.

C. Use Amazon Rekognition for image processing.

D. Use Amazon CloudFront in front of ImageBucket.

E. Deploy the applications in an Amazon ECS cluster and apply Service Auto Scaling.

ANSWER: B D

Explanation:

Use AWS Lambda to run the image processing tasks because image uploads can be handled as an event-driven workflow. Amazon S3 can generate an event notification when a user uploads an object, invoking a Lambda function that runs the required Python image-processing libraries and writes the processed result to ImageBucket. Lambda automatically scales concurrency as upload volume grows, while removing the need to provision, patch, monitor, and scale individual EC2 instances. The Python dependencies can be delivered in the function deployment package, a Lambda layer, or a Lambda container image. This approach provides elastic processing capacity with substantially lower operational management. AWS describes the integration and event-notification pattern in Using AWS Lambda with Amazon S3.

Use Amazon CloudFront in front of ImageBucket because the processed photos are static content consumed by globally distributed users. CloudFront caches images at edge locations close to viewers, reducing latency and improving load performance while minimizing repeated retrieval requests to the S3 origin. CloudFront also scales automatically to accommodate globally distributed read traffic, so the company does not need to operate caching infrastructure. ImageBucket should be configured as a protected CloudFront origin, preferably using Origin Access Control, so users retrieve the images through the distribution. AWS documents S3-origin delivery through CloudFront at Use an Amazon S3 bucket as the origin.

QUESTION NO: 26

A solutions architect must implement a multi-Region architecture for an Amazon RDS for PostgreSQL database that supports a web application. The database launches from an AWS CloudFormation template that includes AWS services and features that are present in both the primary and secondary Regions. The database is configured for automated backups, and it has an RTO of 15 minutes and an RPO of 2 hours. The web application is configured to use an Amazon Route 53 record to route traffic to the database. Which combination of steps will result in a highly available architecture that meets all the requirements? (Choose two.)

A. Create a cross-Region read replica of the database in the secondary Region. Configure an AWS Lambda function in the secondary Region to promote the read replica during a failover event.

B. In the primary Region, create a health check on the database that will invoke an AWS Lambda function when a failure is detected. Program the Lambda function to recreate the database from the latest database snapshot in the secondary Region and update the Route 53 host records for the database.

C. Create an AWS Lambda function to copy the latest automated backup to the secondary Region every 2 hours.

D. Create a failover routing policy in Route 53 for the database DNS record. Set the primary and secondary endpoints to the endpoints in each Region.

E. Create a hot standby database in the secondary Region. Use an AWS Lambda function to restore the secondary database to the latest RDS automatic backup in the event that the primary database fails.

ANSWER: A D

Explanation:

Creating a cross-Region read replica in the secondary Region provides a continuously available, asynchronously replicated copy of the Amazon RDS for PostgreSQL database outside the primary Region. Amazon RDS supports cross-Region read replicas for PostgreSQL, and the replica can be promoted to a standalone, writable DB instance when the primary Region has a failure. Automating promotion through an AWS Lambda function reduces manual intervention and supports the required 15-minute recovery time objective. Because replication is ongoing rather than based on periodic snapshot restores, this approach provides recovery-point protection well within the stated two-hour RPO under normal replication conditions. AWS documents cross-Region replication and replica promotion at Working with cross-Region read replicas.

Creating a Route 53 failover routing policy for the database DNS record supplies the required endpoint redirection mechanism. The primary record directs application connections to the primary Regional database endpoint while it is considered healthy. After the secondary replica is promoted, the secondary failover record can direct DNS responses to the promoted database endpoint, allowing the application to reconnect to the recovered database in the secondary Region. Route 53 failover routing is specifically designed to return a secondary record when the primary record is unhealthy, as described in the Amazon Route 53 failover routing documentation.

QUESTION NO: 27

A company is planning to migrate its on-premises data analysis application to AWS. The application is hosted across a fleet of servers and requires consistent system time.

The company has established an AWS Direct Connect connection from its on-premises data center to AWS. The company has a high-precision stratum-0 atomic clock network appliance that acts as an NTP source for all on-premises servers.

After the migration to AWS is complete, the clock on all Amazon EC2 instances that host the application must be synchronized with the on-premises atomic clock network appliance.

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

A. Configure a DHCP options set with the on-premises NTP server address. Assign the options set to the VPC.

B. Create a custom AMI to use the Amazon Time Sync Service at 169.254.169.123. Use this AMI for the application. Use AWS Config to audit the NTP configuration.

C. Deploy a third-party time server from the AWS Marketplace. Configure the time server to synchronize with the on-premises atomic clock network appliance. Ensure that NTP traffic is allowed inbound in the network ACLs for the VPC that contains the third-party server.

D. Create an IPsec VPN tunnel from the on-premises atomic clock network appliance to the VPC to encrypt the traffic over the Direct Connect connection. Configure the VPC route tables to direct NTP traffic over the tunnel.

ANSWER: A

Explanation:

Configuring a DHCP options set with the on-premises NTP server address and associating it with the VPC is the lowest-administration solution because Amazon VPC DHCP option sets centrally distribute network configuration to instances in that VPC. The `ntp-servers` DHCP option supports specifying NTP server IP addresses or domain names, allowing the company to advertise the existing on-premises atomic-clock appliance as the authoritative time source for the migrated EC2 fleet. This retains the stated requirement that EC2 instance clocks synchronize specifically with the on-premises appliance, rather than using a separate AWS time source.

With Direct Connect already established, EC2 instances can reach the private on-premises NTP endpoint through the hybrid network connection, provided that the required routes and UDP port 123 security controls are in place. Associating one DHCP options set with the VPC avoids maintaining a custom AMI, deploying and operating intermediary time servers, or configuring individual instances. New instances launched in the VPC receive the NTP configuration through DHCP, making the approach scalable as the application fleet changes. AWS documents NTP server configuration as a supported DHCP option and documents that a DHCP options set can be associated with a VPC. See DHCP option sets in Amazon VPC and DHCP option set concepts.

QUESTION NO: 28

A team collects and routes behavioral data for an entire company The company runs a Multi-AZ VPC environment with public subnets, private subnets, and in internet gateway Each public subnet also contains a NAT gateway Most of the company ' s applications read from and write to Amazon Kinesis Data Streams. Most of the workloads am in private subnets.

A solutions architect must review the infrastructure The solutions architect needs to reduce costs and maintain the function of the applications The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.

What should the solutions architect do to meet these requirements?

A. Enable VPC Flow Logs. Use Amazon Athena to analyze the logs for traffic that can be removed. Ensure that security groups are Mocking traffic that is responsible for high costs.

B. Add an interface VPC endpoint for Kinesis Data Streams to the VP

C. Enable VPC Flow Logs and Amazon Detective Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic

D. Add an interface VPC endpoint for Kinesis Data Streams to the VP

ANSWER: D

Explanation:

Add an interface VPC endpoint for Kinesis Data Streams to the VP is the appropriate solution because it provides private connectivity from workloads in the VPC to the Kinesis Data Streams service through AWS PrivateLink. Currently, instances in private subnets likely reach the public Kinesis service endpoint by routing traffic through their NAT gateways. NAT gateways incur an hourly charge and a data-processing charge for every gigabyte processed, so high-volume producers and consumers of streaming data can create substantial `NatGateway-Bytes` costs.

An interface endpoint creates elastic network interfaces with private IP addresses in selected subnets. With private DNS enabled, the standard Kinesis Data Streams endpoint name resolves to these private endpoint addresses for resources in the VPC. Application calls can therefore remain on AWS private networking and no longer require NAT gateway traversal for Kinesis API and data-plane traffic. This maintains the applications' existing purpose while directly addressing the identified source of cost. Endpoint security groups and endpoint policies can also be used to restrict which workloads and Kinesis actions are permitted. AWS documents Kinesis Data Streams support for interface endpoints at Using Amazon Kinesis Data Streams with interface VPC endpoints and explains the PrivateLink interface endpoint model at Access an AWS service using an interface VPC endpoint.

QUESTION NO: 29

A company needs to improve the security of its web-based application on AWS. The application uses Amazon CloudFront with two custom origins. The first custom origin routes requests to an Amazon API Gateway HTTP API. The second custom origin routes traffic to an Application Load Balancer (ALB) The application integrates with an OpenID Connect (OIDC) identity provider (IdP) for user management.

A security audit shows that a JSON Web Token (JWT) authorizer provides access to the API The security audit also shows that the ALB accepts requests from unauthenticated users

A solutions architect must design a solution to ensure that all backend services respond to only authenticated users

Which solution will meet this requirement?

A. Configure the ALB to enforce authentication and authorization by integrating the ALB with the IdP Allow only authenticated users to access the backend services

B. Modify the CloudFront configuration to use signed URLs Implement a permissive signing policy that allows any request to access the backend services

C. Create an AWS WAF web ACL that filters out unauthenticated requests at the ALB level. Allow only authenticated traffic to reach the backend services.

D. Enable AWS CloudTrail to log all requests that come to the ALB Create an AWS Lambda function to analyze the togs and block any requests that come from unauthenticated users.

ANSWER: A

Explanation:

Configure the ALB to enforce authentication and authorization by integrating the ALB with the IdP. This completes authentication enforcement across both application paths. The API Gateway HTTP API already uses a JWT authorizer, which validates JWTs issued by a supported OIDC or OAuth 2.0 identity provider before API requests are authorized. The remaining exposure is the workload reached through the ALB.

An Application Load Balancer listener rule can use an `authenticate-oidc` action before its forwarding action. When a user does not have a valid authenticated session, the ALB redirects the user to the configured OIDC IdP. After the IdP successfully authenticates the user, the ALB validates the authentication response, establishes its authentication session, and forwards the request to the target group. Consequently, backend targets receive requests only after the ALB authentication step has succeeded. This uses managed, native ALB authentication rather than requiring application changes or a custom request-processing component.

The configuration should protect the applicable listener rules and use HTTPS for the authentication flow. This directly addresses the audit finding while retaining the existing JWT-based authorization control for the HTTP API. See Authenticate users using an Application Load Balancer and Control access to HTTP APIs with JWT authorizers.

QUESTION NO: 30

An online magazine will launch Its latest edition this month. This edition will be the first to be distributed globally. The magazine's dynamic website currently uses an Application Load Balancer in front of the web tier a fleet of Amazon EC2 instances for web and application servers, and Amazon Aurora MySQL. Portions of the website include static content and almost all traffic is read-only

The magazine is expecting a significant spike m internet traffic when the new edition is launched Optimal performance is a top priority for the week following the launch

Which combination of steps should a solutions architect take to reduce system response antes for a global audience? (Select TWO)

A. Use logical cross-Region replication to replicate the Aurora MySQL database to a secondary Region Replace the web servers with Amazon S3 Deploy S3 buckets in cross-Region replication mode

B. Ensure the web and application tiers are each m Auto Scaling groups. Introduce an AWS Direct Connect connection Deploy the web and application tiers in Regions across the world

C. Migrate the database from Amazon Aurora to Amazon RDS for MySQL. Ensure all three of the application tiers—web. application, and database—are in private subnets.

D. Use an Aurora global database for physical cross-Region replication. Use Amazon S3 with cross-Region replication for static content and resources. Deploy the web and application tiers in Regions across the world

E. Introduce Amazon Route 53 with latency-based routing and Amazon CloudFront distributions. Ensure me web and application tiers are each in Auto Scaling groups

ANSWER: D E

Explanation:

Use an Aurora global database for physical cross-Region replication. Use Amazon S3 with cross-Region replication for static content and resources. Deploy the web and application tiers in Regions across the world is correct because it places both read capacity and application processing nearer to global users. Aurora Global Database is specifically intended for globally distributed applications that need low-latency local reads. Its secondary AWS Regions provide read-only Aurora DB clusters that are physically replicated from the primary Region, an especially strong fit for this predominantly read-only workload. Replicating static resources to regional S3 buckets and running the web and application tiers in multiple Regions further reduces inter-Region and internet round trips for requests served in each geography.

Introduce Amazon Route 53 with latency-based routing and Amazon CloudFront distributions. Ensure me web and application tiers are each in Auto Scaling groups is also correct because Route 53 latency-based routing directs each user toward the regional endpoint with the lowest observed network latency. CloudFront serves cached static and cacheable web content from AWS edge locations close to viewers, minimizing origin fetches during the launch spike. Auto Scaling allows the web and application tiers to add capacity as demand rises, preserving responsiveness while the edition receives unusually high traffic. Together, these services combine edge caching, latency-aware regional routing, and elastic compute capacity. See Aurora Global Database documentation and Route 53 latency-based routing documentation.

QUESTION NO: 31

A company is migrating an on-premises application and a MySQL database to AWS. The application processes highly sensitive data, and new data is constantly updated in the database. The data must not be transferred over the internet. The company also must encrypt the data in transit and at rest.

The database is 5 TB in size. The company already has created the database schema in an Amazon RDS for MySQL DB instance The company has set up a 1 Gbps AWS Direct Connect connection to AWS. The company also has set up a public VIF and a private VIF. A solutions architect needs to design a solution that will migrate the data to AWS with the least possible downtime

Which solution will meet these requirements?

A. Perform a database backup. Copy the backup files to an AWS Snowball Edge Storage Optimized device. Import the backup to Amazon S3. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest Use TLS for encryption in transit Import the data from Amazon S3 to the DB instance.

B. Use AWS Database Migration Service (AWS DMS) to migrate the data to AWS. Create a DMS replication instance in a private subnet. Create VPC endpoints for AWS DMS. Configure a DMS task to copy data from the on-premises database to the DB instance by using full load plus change data capture (CDC). Use the AWS Key Management Service (AWS KMS) default key for encryption at rest. Use TLS for encryption in transit.

C. Perform a database backup. Use AWS DataSync to transfer the backup files to Amazon S3 Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest. Use TLS for encryption in transit Import the data from Amazon S3 to the DB instance.

D. Use Amazon S3 File Gateway Set up a private connection to Amazon S3 by using AWS PrivateLink. Perform a database backup. Copy the backup files to Amazon S3. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest. Use TLS for encryption in transit. Import the data from Amazon S3 to the DB instance.

ANSWER: B

Explanation:

Use AWS Database Migration Service (AWS DMS) to migrate the data to AWS. Create a DMS replication instance in a private subnet. Create VPC endpoints for AWS DMS. Configure a DMS task to copy data from the on-premises database to the DB instance by using full load plus change data capture (CDC). Use the AWS Key Management Service (AWS KMS) default key for encryption at rest. Use TLS for encryption in transit. This approach is designed for migrations where the source database remains active and the cutover window must be minimized. A full-load-and-change-data-capture task initially copies the existing 5 TB MySQL data, then continually captures and applies source changes to the Amazon RDS for MySQL target. Once replication lag is sufficiently low, the application can briefly stop writes, allow the remaining changes to apply, and switch to RDS. The DMS replication instance can reside in private subnets and reach the on-premises MySQL source through the Direct Connect private virtual interface, keeping database migration traffic off the public internet. DMS supports SSL/TLS for encrypted endpoint connections. Its replication instance storage is encrypted at rest with an AWS KMS key, including use of the AWS managed default key where appropriate. AWS documents full load and ongoing replication in its DMS change data capture documentation and describes private VIF connectivity in the AWS Direct Connect virtual interface guide.

QUESTION NO: 32

A company is using Amazon SageMaker A1 Notebook Instances and SageMaker APIs to train machine learning (ML) models. The SageMaker A1 Notebook Instances are deployed in a VPC that does not have access to or from the internet.

Datasets for ML model training are stored in an Amazon S3 bucket. Interface VPC endpoints provide access to Amazon S3 and the SageMaker APIs.

Occasionally, data scientists require access to a private Git repository to update application packages that they use as part of their workflow. The company must provide access to the Git repository while ensuring that the SageMaker A1 Notebook Instances remain isolated from the internet.

Which solution meets these requirements with the LEAST operational overhead?

A. Add the Git repository as a resource for SageMaker by referencing the remote URL. Configure AWS Secrets Manager to use Git credentials to access the repository.

B. Add the Git repository as a resource for SageMaker by referencing the remote URL. Add the username to the URL that is required to access the repository.

C. Create a NAT gateway in the VP

D. Create a NAT gateway in the VP

E. Migrate or mirror the private Git repository to AWS CodeCommit, and create interface VPC endpoints for CodeCommit API and Git operations.

ANSWER: E

Explanation:

Migrate or mirror the private Git repository to AWS CodeCommit, and create interface VPC endpoints for both the CodeCommit API and CodeCommit Git operations. This provides a fully private path for Git clone, pull, and push traffic from the SageMaker notebook instances without adding an internet gateway or NAT gateway. The Git operations endpoint enables HTTPS Git traffic to CodeCommit through AWS PrivateLink, while the CodeCommit API endpoint supports CodeCommit service API calls. Private DNS can resolve the standard CodeCommit service names to the endpoint network interfaces within the VPC, so notebook instances do not require public addressing or internet egress. Access can be controlled through the notebook instance IAM role and the CodeCommit repository policy, avoiding separately maintained Git credentials and reducing operational administration. SageMaker can associate CodeCommit repositories with notebook instances, allowing the repository to be available as part of the managed notebook workflow. This combines native SageMaker repository integration with private AWS networking and managed IAM authorization. See AWS guidance on interface VPC endpoints for CodeCommit and using Git repositories with SageMaker notebook instances.

QUESTION NO: 33

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:

“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.” permits the application instances to initiate connections only to network interfaces associated with the Aurora DB cluster security group and only on the database listener port. This narrows egress access to the required database service rather than allowing broad outbound network traffic.

“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.” allows connections to the Aurora cluster only from EC2 instances that have the designated application security group attached. Security-group references provide a robust least-privilege control because access follows group membership; scaling or replacing application instances does not require updating IP-based rules. The database port must match the Aurora engine and configured DB cluster port, such as 3306 for Aurora MySQL-compatible editions or 5432 for Aurora PostgreSQL-compatible editions. Security groups are stateful, so response traffic for an allowed connection is automatically allowed without separately opening ephemeral return ports. AWS documents security-group referencing and stateful connection tracking in the Amazon VPC security group rules documentation and explains DB access controls in the Amazon Aurora security groups documentation.

QUESTION NO: 34

A company runs a simple Linux application on Amazon EKS by using nodes of the M6i (general purpose) instance type. The company has an EC2 Instance Savings Plan for the M6i family that will expire soon.

A solutions architect must minimize the EKS compute costs when the Savings Plan expires.

Which combination of steps will meet this requirement? (Select THREE.)

A. Rebuild the application container images to support ARM64 architecture.

B. Rebuild the application container images to support containers.

C. Migrate the EKS nodes to the most recent generation of Graviton-based instances.

D. Replace the EKS nodes with the most recent generation of x86_64 instances.

E. Purchase a new EC2 Instance Savings Plan for the newly selected Graviton instance family.

F. Purchase a new EC2 Instance Savings Plan for the newly selected x86_64 instance family.

ANSWER: A C E

Explanation:

Rebuild the application container images to support ARM64 architecture, migrate the EKS nodes to the most recent generation of Graviton-based instances, and purchase a new EC2 Instance Savings Plan for the newly selected Graviton instance family provide a cost-optimized migration path. AWS Graviton processors use the Arm64 instruction set and generally offer better price performance than comparable x86-based instances. The application’s container images therefore must include an ARM64-compatible build before pods can run on Graviton worker nodes. For a simple Linux application, this commonly means rebuilding the image for ARM64 or publishing a multi-architecture image that supports both amd64 and arm64 during the transition.

After confirming the workload runs correctly on Arm64, moving the EKS node group to a current Graviton general-purpose instance family, such as M7g or a newer compatible generation, reduces the On-Demand compute rate while retaining the general-purpose profile appropriate for the existing M6i-based nodes. An EC2 Instance Savings Plan for that selected instance family and AWS Region then provides additional discounts in exchange for a one- or three-year hourly spending commitment. The commitment should be sized from the expected steady-state Graviton node usage. AWS documents EKS Arm64 scheduling and image requirements at Amazon EKS Arm support and describes instance-family flexibility and commitments at Savings Plans.

QUESTION NO: 35

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 appropriate troubleshooting action for an insufficient-capacity error when adding instances to a placement group, particularly a cluster placement group. Cluster placement groups aim to locate instances physically close together within one Availability Zone to provide very low network latency and high network throughput. As capacity in the required physical area becomes fragmented, EC2 might be unable to place an additional instance close enough to the already running instances, even when the Availability Zone has general instance capacity available.

Stopping all instances in the placement group and then starting them again allows Amazon EC2 to attempt a new placement for the group. This can enable EC2 to find suitable contiguous capacity for the existing workload and the newly requested instances. The instances must be stopped rather than rebooted because a reboot retains the existing host placement; a stop/start cycle can move EBS-backed instances to different underlying hardware. After the placement group has been restarted, retry the instance launch. AWS explicitly recommends this procedure when there is insufficient capacity to launch an instance into a cluster placement group. See the Amazon EC2 placement group strategies documentation and the EC2 stop and start documentation.

QUESTION NO: 36

A company hosts a game player-matching service on a public-facing, physical, on-premises instance that all users are able to access over the instance uses UDP. The company wants to migrate the service to AWS and provide a high level of security. A solutions architect needs to de matching service using AWS.

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

E.

)

A. Use a Network Load Balancer (NLB) in front of the player-matching instance. Use a friendly DNS entry in Amazon Route 53 that points to the NLB.

B. Use an Application Load Balancer (ALB) in front of the player-matching instance. Use a friendly DNS entry in Amazon Route 53 that points to the public-facing fully qualified domain name (FQDN).

C. Define an AWS WAF rule to explicitly drop non-UDP traffic, and associate the rule with the load balancer.

D. Configure a network ACL rule to block all non-UDP traffic. Associate the network ACL with the subnets that hold the load balancer.

E. Configure AWS Shield Advanced protection for the Network Load Balancer.

ANSWER: A D E

Explanation:

Use a Network Load Balancer (NLB) in front of the player-matching instance and a friendly Amazon Route 53 DNS record that points to the NLB because an NLB operates at Layer 4 and supports UDP listeners. This allows the UDP-based game

player-matching service to be exposed through a scalable AWS-managed entry point while presenting users with a stable, friendly DNS name. A network ACL associated with the load balancer subnets can enforce subnet-level protocol filtering by allowing the required UDP traffic and blocking unwanted non-UDP traffic. Because network ACLs are stateless, the rule set must also allow the necessary return traffic and ephemeral port ranges for the application's traffic flow. Configure AWS Shield Advanced protection for the Network Load Balancer to add enhanced protection against DDoS attacks for this public-facing service. Shield Advanced supports Network Load Balancers and provides additional detection, mitigation, and visibility capabilities for protected resources. Together, the NLB, network ACL controls, and Shield Advanced provide UDP support, controlled network access, and stronger resilience against volumetric and application-impacting DDoS events. See Network Load Balancer documentation and AWS Shield Advanced protections.

QUESTION NO: 37

A company hosts a metadata API on Amazon EC2 instances behind an internet-facing Application Load Balancer (ALB). Only internal applications that run on EC2 instances in separate AWS accounts need to access the metadata API. All the internal EC2 instances use NAT gateways.

A new policy requires that traffic between internal applications must not travel across the public internet.

Which solution will meet this requirement?

A. Create an HTTP API in Amazon API Gateway. Configure a route for the metadata API. Configure a VPC link to the VPC that hosts the metadata API's EC2 instances. Update the API Gateway resource policy to include the account IDs of the internal applications that access the metadata API.

B. Create a REST API in Amazon API Gateway. Specify the API Gateway endpoint type as private. Associate the REST API with the metadata API's VPC. Create a gateway VPC endpoint for the REST API. Share the endpoint across accounts by using AWS Resource Access Manager (AWS RAM). Configure the internal applications to connect to the gateway VPC endpoint.

C. Create an internal ALB. Register the metadata API's EC2 instances with the internal ALB. Create an internal Network Load Balancer (NLB) that has a target group type of ALB. Register the internal ALB as the target. Configure an AWS PrivateLink endpoint service for the NLB. Grant the internal applications access to the metadata API through the PrivateLink endpoint.

D. Create an internal ALB. Register the metadata API's EC2 instances with the internal ALB. Configure an AWS PrivateLink endpoint service for the internal ALB. Grant the internal applications access to the metadata API through the PrivateLink endpoint.

ANSWER: C

Explanation:

Create an internal ALB, register the metadata API EC2 instances as its targets, and place an internal Network Load Balancer in front of that ALB by using an ALB-type target group. The Network Load Balancer can then be published as an AWS PrivateLink endpoint service. PrivateLink allows consumers in the separate AWS accounts to create interface endpoints in their own VPCs and connect to the service through private IP addresses. This keeps the application-to-service path on the AWS network rather than routing requests through NAT gateways to the public address of the existing internet-facing load balancer.

The internal ALB continues to provide layer 7 HTTP routing and health checks for the API targets. The NLB supplies the endpoint-service integration required by PrivateLink, because endpoint services are backed by Network Load Balancers (or Gateway Load Balancers), not directly by an Application Load Balancer. An NLB target group with target type `alb` is specifically supported for forwarding NLB traffic to an ALB. The service provider can grant access to the intended AWS principals, and each consumer account can use its approved interface endpoint privately and without requiring VPC peering or public internet connectivity. See AWS PrivateLink endpoint services and ALB as a target of an NLB.

QUESTION NO: 38

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:

Creating an SCP that uses the `aws:RequestedRegion` global condition key to deny requests outside `ap-northeast-1`, and attaching it to the root, provides an organization-wide preventative guardrail. SCPs define the maximum available permissions for member accounts, and a policy attached at the root is inherited by every OU and account beneath it. This centrally enforces the Regional requirement for current and future accounts without requiring individual IAM users, roles, or account-level policies to be maintained. AWS documents this approach as the standard SCP pattern for restricting access to approved AWS Regions. See AWS Organizations SCP examples.

Creating an SCP with the `ec2:InstanceType` condition key and applying it to the DataOps OU provides a targeted preventative control for all accounts in that OU. An SCP can deny EC2 instance launch requests when the requested instance type is not in the approved list. Because the policy is inherited by accounts under DataOps, the restriction remains effective as accounts are added or reorganized, while Research is unaffected by this DataOps-specific requirement. The `ec2:InstanceType` condition key is supported by applicable EC2 authorization actions, including instance launch controls. See the Amazon EC2 service authorization reference.

QUESTION NO: 39

A company is using AWS CloudFormation as its deployment tool for all applications. It stages all application binaries and templates within Amazon S3 buckets with versioning enabled. Developers have access to an Amazon EC2 instance that hosts the integrated development environment (IDE). The developers download the application binaries from Amazon S3 to the EC2 instance, make changes, and upload the binaries to an S3 bucket after running the unit tests locally. The developers want to improve the existing deployment mechanism and implement CI/CD using AWS CodePipeline. The developers have the following requirements: • Use AWS CodeCommit for source control. • Automate unit testing and security scanning. • Alert the developers when unit tests fail. • Turn application features on and off, and customize deployment dynamically as part of CI/CD. • Have the lead developer provide approval before deploying an application. Which solution will meet these requirements?

A. Use AWS CodeBuild to run unit tests and security scans. Use an Amazon EventBridge rule to send Amazon SNS alerts to the developers when unit tests fail. Write AWS Cloud Development Kit (AWS CDK) constructs for different solution features, and use a manifest file to tum features on and off in the AWS CDK application. Use a manual approval stage in the pipeline to allow the lead developer to approve applications.

B. Use AWS Lambda to run unit tests and security scans. Use Lambda in a subsequent stage in the pipeline to send Amazon SNS alerts to the developers when unit tests fail. Write AWS Amplify plugins for different solution features and utilize user prompts to tum features on and off. Use Amazon SES in the pipeline to allow the lead developer to approve applications.

C. Use Jenkins to run unit tests and security scans. Use an Amazon EventBridge rule in the pipeline to send Amazon SES alerts to the developers when unit tests fail Use AWS CloudFormation nested stacks for different solution features and parameters to turn features on and off. Use AWS Lambda in the pipeline to allow the lead developer to approve applications.

D. Use AWS CodeDeploy to run unit tests and security scans. Use an Amazon CloudWatch alarm in the pipeline to send Amazon SNS alerts to the developers when unit tests fail. Use Docker images for different solution features and the AWS CLI to turn features on and off. Use a manual approval stage in the pipeline to allow the lead developer to approve applications.

ANSWER: A

Explanation:

Use AWS CodeBuild to run unit tests and security scans. Use an Amazon EventBridge rule to send Amazon SNS alerts to the developers when unit tests fail. Write AWS Cloud Development Kit (AWS CDK) constructs for different solution features, and use a manifest file to tum features on and off in the AWS CDK application. Use a manual approval stage in the pipeline to allow the lead developer to approve applications. provides an AWS-native CI/CD architecture that satisfies the stated workflow. AWS CodePipeline can use AWS CodeCommit as its source provider, enabling commits to initiate a pipeline execution. A CodeBuild action can run a repeatable build specification containing the required unit-test commands and security-scanning tools. Build failures result in pipeline and CodeBuild state-change events, which Amazon EventBridge can route to an Amazon SNS topic so the development team receives notifications without adding custom notification code. AWS CDK allows infrastructure definitions to use programming-language constructs and configuration inputs, including a manifest, to conditionally include resources and tailor the CloudFormation templates synthesized for a deployment. This makes feature enablement and deployment customization part of the controlled pipeline process. Finally, a CodePipeline manual approval action pauses progression before deployment until the designated lead developer approves or rejects the release. See AWS CodePipeline manual approval actions and AWS CodeBuild documentation.

QUESTION NO: 40

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 ofthe 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:

Cross-account retrieval of objects encrypted with a customer managed AWS KMS key requires authorization from both Amazon S3 and AWS KMS. “Create a bucket policy that includes read permissions for the S3 bucket. Set the principal ofthe bucket policy to the account ID of the Strategy account” establishes the S3 resource-based permission that allows the external account to access the required bucket objects. This read access should be scoped to the needed actions, such as `s3:GetObject`, and relevant object prefixes where appropriate.

“Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key” gives the assumed role the identity-based permissions necessary to retrieve the objects and request decryption. “Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role” is also necessary because a customer managed KMS key must explicitly authorize cross-account use. The KMS permission should be limited to `kms:Decrypt` for the specific key, optionally constrained so it is used through Amazon S3. Together, these permissions let strategy users view the encrypted content while preserving least privilege. See Amazon S3 bucket policy examples and AWS KMS cross-account key access guidance.

QUESTION NO: 41

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 A.

B. In Account A, set the S3 bucket policy to the following:

C. In Account A, set an S3 bucket policy that allows Account B's User_DataProcessor principal to perform the required actions, such as s3:GetObject on the shared objects.

D. In Account B, attach an IAM policy to User_DataProcessor that allows the required S3 actions, such as s3:GetObject on the objects in Account A's bucket.

E. In Account B, set the permissions of User_DataProcessor to the following:

ANSWER: C D

Explanation:

Cross-account Amazon S3 access requires authorization from both the account that owns the bucket and the account that owns the requesting IAM user. In Account A, the S3 bucket policy must explicitly allow the external principal, such as the `User_DataProcessor` IAM user ARN or an appropriately scoped principal in Account B, to perform the required actions. For downloading the data files, this normally includes `s3:GetObject` against the object ARN pattern. If the user must browse or enumerate files, the policy should also allow `s3:ListBucket` against the bucket ARN.

In Account B, `User_DataProcessor` must have an identity-based IAM policy that permits the same required S3 actions for Account A’s bucket and its objects. This ensures that the user is authorized by its own account while the bucket owner retains control over external access through its bucket policy. AWS evaluates the applicable identity-based and resource-based policies for the request; therefore, both sides of this cross-account permission model must be configured with the intended least-privilege scope. AWS provides this cross-account S3 pattern in its cross-account bucket permissions walkthrough and explains the relevant S3 actions and resource formats in the Amazon S3 policy actions documentation.

QUESTION NO: 42

A company is developing a new service that will be accessed using TCP on a static port A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name myservice.com, which is publicly accessible The service must use fixed address assignments so other companies can add the addresses to their allow lists.

Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?

A. Create Amazon EC2 instances with an Elastic IP address for each instance Create a Network Load Balancer (NLB) and expose the static TCP port Register EC2instances with the NLB Create a new name server record set named my service com, and assign the Elastic IP addresses of the EC2 instances to the record set Provide the Elastic IP addresses of the EC2 instances to the other companies to add to their allow lists

B. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP addresses for the ECS cluster Create a Network Load Balancer (NLB) and expose the TCP port Create a target group and assign the ECS cluster name to the NLB Create a new A record set named my service com and assign the public IP addresses of the ECS cluster to the record set Provide the public IP addresses of the ECS cluster to the other com

C. Create Amazon EC2 instances for the service Create one Elastic IP address for each Availability Zone Create a Network Load Balancer (NLB) and expose the assigned TCP port Assign the Elastic IP addresses to the NLB for each Availability Zone Create a target group and register the EC2 instances with the NLB Create a new A (alias) record set named my service com, and assign the NLB DNS name to the record set.

D. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP address for each host in the cluster Create an Application Load Balancer (ALB) and expose the static TCP port Create a target group and assign the ECS service definition name to the ALB Create a new CNAME record set and associate the public IP addresses to the record set Provide the Elastic IP addresses of the Amazon EC2 instances to the ot

ANSWER: C

Explanation:

Create Amazon EC2 instances for the service Create one Elastic IP address for each Availability Zone Create a Network Load Balancer (NLB) and expose the assigned TCP port Assign the Elastic IP addresses to the NLB for each Availability Zone Create a target group and register the EC2 instances with the NLB Create a new A (alias) record set named my service com, and assign the NLB DNS name to the record set. This design directly meets the TCP, static-address, DNS, and multi-Availability Zone availability requirements. A Network Load Balancer operates at Layer 4 and supports TCP listeners on a specified port. When an internet-facing NLB is enabled in multiple Availability Zones, an Elastic IP address can be assigned to the load balancer's subnet mapping in each enabled Availability Zone. These addresses are static and can therefore be provided to external organizations for IP allow lists. The NLB distributes connections only to healthy registered targets, while deployment across multiple Availability Zones provides resilient load balancer entry points and target placement. A Route 53 A alias record can point the public DNS name to the NLB without requiring the company to manage changing load balancer IP addresses in DNS. The NLB remains the stable service endpoint, and its per-Availability Zone Elastic IP addresses provide the fixed addresses required by partners. See Network Load Balancer documentation and Route 53 routing to Elastic Load Balancing.

QUESTION NO: 43

A company has loT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.

A solutions architect designs a solution in which the loT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically encountering a ReadProvisionedThroughputExceeded error.

Which actions should the solutions architect take to resolve this issue? (Select THREE.)

A. Reshard the stream to increase the number of shards in the stream.

B. Use the Kinesis Producer Library (KPL). Adjust the polling frequency.

C. Use consumers with the enhanced fan-out feature.

D. Reshard the stream to reduce the number of shards in the stream.

E. Use an error retry and exponential backoff mechanism in the consumer logic.

F. Configure the stream to use dynamic partitioning.

ANSWER: A C E

Explanation:

"Reshard the stream to increase the number of shards in the stream" increases the stream's aggregate read capacity. In provisioned mode, each shard supports up to 2 MB/second of read throughput, shared by standard consumers. Adding shards can therefore provide more total capacity, provided records and consumers are distributed across the shards appropriately. "Use consumers with the enhanced fan-out feature" is particularly effective when multiple applications concurrently read the same stream. Enhanced fan-out gives every registered consumer dedicated throughput of up to 2 MB/second per shard, eliminating competition with other enhanced fan-out consumers for shared standard read throughput. "Use an error retry and exponential backoff mechanism in the consumer logic" ensures consumers handle transient throttling safely. Kinesis Data Streams guidance recommends backing off after a `ReadProvisionedThroughputExceeded` exception, which avoids immediate repeated requests and allows the consumer to resume reads reliably when capacity is available. Together, these measures increase available read capacity, isolate concurrent consumers where needed, and make read processing resilient to throttling. See the AWS guidance on troubleshooting Kinesis Data Streams consumers and enhanced fan-out.

QUESTION NO: 44

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 approach is to use VPC sharing, which is implemented through AWS Resource Access Manager (AWS RAM). Enabling resource sharing from the AWS Organizations management account authorizes AWS RAM integration with the organization. This lets the infrastructure account share supported resources directly with organizational units or member accounts, avoiding individual account invitations and supporting centralized administration at organizational scale.

Creating a resource share in AWS Resource Access Manager in the infrastructure account, selecting the organizational unit that needs access, and associating the required subnets implements the shared-network model. The infrastructure account remains the VPC owner and retains control of network components, including the VPC, subnets, route tables, network ACLs, and gateways. Accounts that receive the subnet share become VPC participants. They can create supported resources, such as Amazon EC2 instances and certain other service resources, in the shared subnets while not gaining authority to alter the centrally managed network configuration.

This model precisely separates network governance from workload deployment: the infrastructure team operates the common network, and workload-owning accounts deploy into assigned subnet capacity. AWS documents the ownership and participant responsibilities for this design in VPC sharing and explains organization-wide sharing prerequisites in AWS RAM resource sharing.

QUESTION NO: 45

A company is planning to migrate 1,000 on-premises servers to AWS. The servers run on several VMware clusters in the company's data center. As part of the migration plan, the company wants to gather server metrics such as CPU details, RAM usage, operating system information, and running processes. The company then wants to query and analyze the data.

Which solution will meet these requirements?

A. Deploy and configure the AWS Agentless Discovery Connector virtual appliance on the on-premises hosts. Configure Data Exploration in AWS Migration Hub. Use AWS Glue to perform an ETL job against the data. Query the data by using Amazon S3 Select.

B. Export only the VM performance information from the on-premises hosts. Directly import the required data into AWS Migration Hub. Update any missing information in Migration Hub. Query the data by using Amazon QuickSight.

C. Create a script to automatically gather the server information from the on-premises hosts. Use the AWS CLI to run the put-resource-attributes command to store the detailed server data in AWS Migration Hub. Query the data directly in the Migration Hub console.

D. Deploy the AWS Application Discovery Agent to each on-premises server. Configure Data Exploration in AWS Migration Hub. Use Amazon Athena to run predefined queries against the data in Amazon S3.

ANSWER: D

Explanation:

Deploy the AWS Application Discovery Agent to each on-premises server. Configure Data Exploration in AWS Migration Hub. Use Amazon Athena to run predefined queries against the data in Amazon S3. This solution provides the required depth of server-level discovery for migration planning. The AWS Application Discovery Agent runs on each server and collects detailed system configuration and utilization data, including CPU and memory information, operating system details, network connections, and running processes. Agent-based collection is particularly appropriate here because the requirement includes process-level information, which must be observed from within the guest operating system rather than solely from VMware inventory data.

AWS Application Discovery Service integrates with AWS Migration Hub to organize the collected inventory and dependency information. Migration Hub Data Exploration makes the discovery dataset available for analysis, including through exported data in Amazon S3. Amazon Athena can query data stored in S3 directly using SQL, without loading it into a separate database. AWS supplies predefined Athena queries for discovery data, enabling the company to analyze server inventory, utilization, and application-related information at the scale needed to plan a migration of 1,000 servers. This combines comprehensive data collection with serverless, query-based analysis. See the AWS Application Discovery Agent documentation and the Amazon Athena User Guide.

QUESTION NO: 46

A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts. Each VPC consists to public subnets and private subnets that span across multiple Availability Zones. NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets.

A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC. The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account.

Which set of additional steps should the solutions architect take to meet these requirements?

A. Create peering connections between the egress VPC and the spoke VPCs. Configure the required routing to allow access to the internet.

B. Create a transit gateway, and share it with the existing AWS accounts. Attach existing VPCs to the transit gateway Configure the required routing to allow access to the internet.

C. Create a transit gateway in every account. Attach the NAT gateway to the transit gateways. Configure the required routing to allow access to the internet.

D. Create an AWS PrivateLink connection between the egress VPC and the spoke VPCs. Configure the required routing to allow access to the internet

ANSWER: B

Explanation:

Create a transit gateway, and share it with the existing AWS accounts. Attach existing VPCs to the transit gateway Configure the required routing to allow access to the internet. This creates the scalable hub required for connecting VPCs across hundreds of accounts while centralizing outbound internet egress. The central account can share the transit gateway with spoke accounts by using AWS Resource Access Manager (AWS RAM), after which each account creates a VPC attachment for its spoke VPC. Private-subnet route tables in each spoke VPC need a default route for internet-bound traffic to the transit gateway. The transit gateway route table then directs that traffic from spoke attachments to the egress VPC attachment. Within the egress VPC, the transit gateway attachment subnet routes the traffic to the NAT gateway, and the NAT gateway's public subnet has a route to an internet gateway. Return routing must also allow traffic from the egress VPC and transit gateway back to the appropriate spoke attachments. For high availability and to avoid unnecessary cross-AZ data transfer, the design should use transit gateway attachment subnets and NAT gateways in multiple Availability Zones, with AZ-aligned routing. AWS documents this as a centralized egress pattern using AWS Transit Gateway and NAT gateways. See Centralized NAT gateway egress with Transit Gateway and Sharing a transit gateway.

QUESTION NO: 47

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 DynamoDB. Use Amazon Simple Notification Service (Amazon SNS) for asynchronous messaging.

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. Use Amazon Simple Queue Service (Amazon SQS) for asynchronous messaging. This design aligns with the application’s container, infrastructure-control, relational-data, and asynchronous-communication requirements. Amazon ECS provides managed orchestration for containerized applications, including scheduling, service discovery integration, deployments, and scaling. Using the Amazon EC2 launch type means the company chooses and manages the EC2 capacity that runs container workloads. This enables the required control over the host operating system, instance configuration, and network configuration while retaining ECS orchestration capabilities. AWS documents the EC2 launch type and its capacity management model at Amazon ECS launch types.

Amazon Aurora MySQL Serverless v2 supplies a managed, strongly relational database with SQL capabilities, transactions, relationships, and automatic capacity scaling. It is therefore suitable for a refactored application whose domain and

persistence model remain relational. Amazon SQS provides durable, managed queues that decouple .NET Core services and support asynchronous producer-consumer processing, making it an appropriate AWS-native replacement pattern for queue-based MSMQ communication. See What is Amazon SQS? for its core messaging model and delivery characteristics.

QUESTION NO: 48

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 m 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:

“Deploy the security tool on EC2 instances m a new Auto Scaling group in the existing VPC” provides a scalable and highly available fleet for the existing third-party appliance. Because the vendor does not offer a cloud-native version, hosting the tool on EC2 preserves the required inspection capability. An Auto Scaling group can distribute appliance instances across multiple Availability Zones and replace unhealthy instances, supporting resilient, real-time processing capacity.

“Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool” provides the purpose-built service insertion mechanism for transparently routing traffic through virtual security appliances. Gateway Load Balancer operates at Layer 3 and uses the GENEVE protocol to encapsulate and send traffic to healthy appliance targets while preserving packet flow information needed by inline inspection tools. Gateway Load Balancer endpoints are deployed in the relevant Availability Zones and VPC route tables direct ingress and egress traffic through the inspection fleet. This architecture lets the application remain independent of the inspection implementation while providing horizontal scale and high availability within the Region. AWS documents Gateway Load Balancer as a service for deploying and scaling third-party virtual appliances, including firewalls, intrusion detection systems, and deep packet inspection solutions. See Gateway Load Balancer documentation and Gateway Load Balancer endpoints for AWS PrivateLink.

QUESTION NO: 49

A company is designing an AWS environment tor a manufacturing application. The application has been successful with customers, and the application ' s user base has increased. The company has connected the AWS environment to the company ' s on-premises data center through a 1 Gbps AWS Direct Connect connection. The company has configured BGP for the connection.

The company must update the existing network connectivity solution to ensure that the solution is highly available, fault tolerant, and secure.

Which solution win meet these requirements MOST cost-effectively?

A. Add a dynamic private IP AWS Site-to-Site VPN as a secondary path to secure data in transit and provide resilience for the Direct Conned connection. Configure MACsec to encrypt traffic inside the Direct Connect connection.

B. Provision another Direct Conned connection between the company ' s on-premises data center and AWS to increase the transfer speed and provide resilience. Configure MACsec to encrypt traffic inside the Dried Conned connection.

C. Configure multiple private VIFs. Load balance data across the VIFs between the on-premises data center and AWS to provide resilience.

D. Add a static AWS Site-to-Site VPN as a secondary path to secure data in transit and to provide resilience for the Direct Connect connection.

ANSWER: A

Explanation:

Add a dynamic private IP AWS Site-to-Site VPN as a secondary path to secure data in transit and provide resilience for the Direct Connect connection. Configure MACsec to encrypt traffic inside the Direct Connect connection. This approach combines an encrypted VPN path with BGP-based dynamic routing, allowing connectivity to fail over automatically when the primary Direct Connect path is unavailable. Dynamic routing is important because route advertisements can be withdrawn and re-advertised during a failure, reducing manual intervention and improving operational resilience. AWS Site-to-Site VPN protects its traffic with IPsec encryption, providing a secure alternate connection path. MACsec provides Layer 2 encryption for traffic traversing an eligible Direct Connect connection, protecting the primary private circuit without requiring the application to implement its own encryption. Together, these controls address availability, fault tolerance, and encryption in transit while avoiding the recurring cost of maintaining an additional dedicated Direct Connect circuit solely as a standby path. AWS recommends using a VPN connection as a backup for Direct Connect and configuring BGP to support automatic failover behavior. See the AWS Direct Connect User Guide and the AWS Site-to-Site VPN User Guide for the supported routing and resilience design patterns.

QUESTION NO: 50

A company has an organization that has many AWS accounts in AWS Organizations. A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.

The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to and from the company's on-premises network.

Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.

The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts.

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

A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS account. Deploy an AWS Lambda function in each AWS account. Configure the Lambda function to run every time an SNS topic receives a message. Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account. Instruct the security team to distribute changes by publishing messages to its SNS topic.

B. Create new customer-managed prefix lists in each AWS account within the organization. Populate the prefix lists in each account with all internal CIDR ranges. Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups. Instruct the security team to share updates with each AWS account owner.

C. Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.

D. Create an IAM role in each account in the organization. Grant permissions to update security groups. Deploy an AWS Lambda function in the security team's AWS account. Configure the Lambda function to take a list of internal IP addresses as input, assume a role in each organization account, and add the list of IP addresses to the security groups in each account.

ANSWER: C

Explanation:

Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups. This provides a centrally owned, reusable representation of the company's approved on-premises network ranges. Instead of copying every CIDR range into security groups throughout every account, account owners reference one shared prefix list in the relevant inbound or outbound rules.

The security team remains the prefix list owner and can add, remove, or modify its entries as network requirements change. Security groups that reference the shared list then use the current entries automatically, eliminating the need to send updates to each account owner or maintain synchronization automation. AWS Resource Access Manager supports sharing customer-managed prefix lists with an AWS Organization or organizational units, making this appropriate for a multi-account environment. The security groups remain managed in their respective accounts, while the common address allow list is governed centrally. This is a native AWS capability and therefore has substantially lower operational overhead than building functions, messaging, and cross-account update workflows. See the AWS documentation for customer-managed prefix lists and AWS RAM shareable resources.

QUESTION NO: 51

A data analytics company has an Amazon Redshift cluster that consists of several reserved nodes. The duster is experiencing unexpected bursts of usage because a team of employees is compiling a deep audit analysis report The queries to generate the report are complex read queries and are CPU intensive.

Business requirements dictate that the cluster must be able to service read and write queries at at) times A solutions architect must devise a solution that accommodates the bursts of usage

Which solution meets these requirements MOST cost-effectively?

A. Provision an Amazon EMR duster Offload the complex data processing tasks

B. Deploy an AWS Lambda function to add capacity to the Amazon Redshift cluster by using a classic resize operation when the duster's CPU metrics in Amazon CloudWatch reach 80%.

C. Deploy an AWS Lambda function to add capacity to the Amazon Redshift duster by using an elastic resize operation when the duster's CPU metrics in Amazon CloudWatch leach 80%.

D. Turn on the Concurrency Scaling feature for the Amazon Redshift duster

ANSWER: D

Explanation:

Turning on Amazon Redshift Concurrency Scaling is the most cost-effective solution because it is designed to absorb intermittent increases in query demand without requiring the company to permanently add provisioned nodes. When workloads exceed the main cluster's available concurrency, Redshift can automatically provision transient concurrency-scaling capacity and route eligible queries to it. This lets the CPU-intensive audit-report reads run with additional capacity while the primary cluster remains available to continue processing its normal workload, including write activity.

The feature is especially suitable for unpredictable or short-lived reporting bursts because capacity is added only when needed and is removed automatically after demand subsides. It therefore avoids the cost and operational disruption of repeatedly resizing the reserved-node cluster to accommodate a temporary peak. Redshift also provides one hour of concurrency-scaling credit per day for each active cluster, accrued up to a limit, which can further reduce the cost of burst capacity. Configuration is managed through Redshift workload management queues, allowing the reporting workload to use concurrency scaling without application-side infrastructure changes. For implementation details and supported workload behavior, see the Amazon Redshift Concurrency Scaling documentation and the Amazon Redshift pricing page.

QUESTION NO: 52

A company ' s solutions architect is analyzing costs of a multi-application environment. The environment is deployed across multiple Availability Zones in a single AWS Region. After a recent acquisition, the company manages two organizations in

AWS Organizations. The company has created multiple service provider applications as AWS PrivateLink-powered VPC endpoint services in one organization. The company has created multiple service consumer applications in the other organization.

Data transfer charges are much higher than the company expected, and the solutions architect needs to reduce the costs. The solutions architect must recommend guidelines for developers to follow when they deploy services. These guidelines must minimize data transfer charges for the whole environment.

Which guidelines meet these requirements? (Select TWO.)

A. Use AWS Resource Access Manager to share the subnets that host the service provider applications with other accounts in the organization.

B. Place the service provider applications and the service consumer applications in AWS accounts in the same organization.

C. Turn off cross-zone load balancing for the Network Load Balancer in all service provider application deployments.

D. Ensure that service consumer compute resources use the Availability Zone-specific endpoint service by using the endpoint ' s local DNS name.

E. Create a Savings Plan that provides adequate coverage for the organization ' s planned inter-Availability Zone data transfer usage.

ANSWER: C D

Explanation:

Turning off cross-zone load balancing for the Network Load Balancer in all service provider application deployments helps retain traffic within the Availability Zone in which it enters the provider service. AWS PrivateLink endpoint services use Network Load Balancers, and disabling cross-zone load balancing prevents the load balancer from sending a request received by a load balancer node in one Availability Zone to a target in another Availability Zone. This reduces chargeable inter-Availability Zone traffic for high-volume service-provider workloads while preserving a multi-AZ deployment model.

Ensuring that service consumer compute resources use the Availability Zone-specific endpoint service by using the endpoint's local DNS name establishes Availability Zone affinity on the consumer side. Interface VPC endpoints create endpoint network interfaces in enabled Availability Zones and expose both Regional and zonal DNS names. Consumers can use the zonal DNS name corresponding to the same Availability Zone as their compute resources, so traffic does not need to cross an Availability Zone boundary to reach an endpoint network interface. Combined with provider-side Network Load Balancer zonal affinity, this approach minimizes unnecessary inter-AZ data transfer across the PrivateLink architecture. AWS documents Network Load Balancer cross-zone behavior at Network Load Balancer cross-zone load balancing and endpoint DNS naming at AWS PrivateLink DNS names.

QUESTION NO: 53

A company is running a large containerized workload in the AWS Cloud. The workload consists of approximately 100 different services. The company uses Amazon ECS to orchestrate the workload.

Recently, the company’s development team started using AWS Fargate instead of Amazon EC2 instances in the ECS cluster. In the past, the workload has come close to running the maximum number of EC2 instances that are available in the account.

The company is worried that the workload could reach the maximum number of ECS tasks that are allowed. A solutions architect must implement a solution that will notify the development team when Fargate reaches 80% of the maximum number of tasks.

What should the solutions architect do to meet this requirement?

A. Use Amazon CloudWatch to monitor the Sample Count statistic for each service in the ECS cluster. Set an alarm for when the math expression sample count/SERVICE_QUOTA service times 100 is greater than 80. Notify the development team by using Amazon SNS.

B. Use Amazon CloudWatch to monitor service quotas that are published under the AWS/Usage metric namespace. Set an alarm for when the math expression metric/SERVICE_QUOTA metric times 100 is greater than 80. Notify the development team by using Amazon SNS.

C. Create an AWS Lambda function to poll detailed metrics from the ECS cluster. When the number of running Fargate tasks is greater than 80, invoke Amazon SES to notify the development team.

D. Create an AWS Config rule to evaluate whether the Fargate service quota is greater than 80. Use Amazon SES to notify the development team when the AWS Config rule is not compliant.

ANSWER: B

Explanation:

Use Amazon CloudWatch to monitor service quotas that are published under the AWS/Usage metric namespace. Set an alarm for when the math expression metric/SERVICE_QUOTA metric times 100 is greater than 80. Notify the development team by using Amazon SNS.

AWS publishes usage metrics for supported service quotas in the `AWS/Usage` namespace, allowing the company to observe its current Fargate task consumption without operating custom polling infrastructure. CloudWatch metric math provides the `SERVICE_QUOTA` function, which retrieves the applicable quota value for a usage metric. Dividing the current usage metric by that quota and multiplying by 100 produces a percentage-based utilization value. An alarm at 80% gives the development team advance notice before the account reaches the task quota and workloads are prevented from launching additional tasks.

Amazon SNS is an appropriate CloudWatch alarm action because it can deliver the notification to subscribed endpoints, such as email addresses, HTTPS integrations, or automated remediation workflows. This approach is managed, quota-aware, and continuously evaluates the actual account-level consumption relative to the current service quota. AWS documents this usage-metric and quota-alarm pattern at Visualize and alarm on service quotas and documents the `SERVICE_QUOTA` metric-math function at Using metric math.

QUESTION NO: 54

A company has deployed applications to thousands of Amazon EC2 instances in an AWS account. A security audit discovers that several unencrypted Amazon Elastic Block Store (Amazon EBS) volumes are attached to the EC2 instances. The company's security policy requires the EBS volumes to be encrypted. The company needs to implement an automated solution to encrypt the EBS volumes. The solution also must prevent development teams from creating unencrypted EBS volumes. Which solution will meet these requirements?

A. Configure the AWS Config managed rule that identifies unencrypted EBS volumes. Configure an automatic remediation action. Associate an AWS Systems Manager Automation runbook that includes the steps to create a new encrypted EBS volume. Create an AWS Key Management Service (AWS KMS) customer managed key. In the key policy, include a statement to deny the creation of unencrypted EBS volumes.

B. Use AWS Systems Manager Fleet Manager to create a list of unencrypted EBS volumes, Create a Systems Manager Automation runbook that includes the steps to create a new encrypted EBS volume. Create an SCP to deny the creation of unencrypted EBS volumes.

C. Use AWS Systems Manager Fleet Manager to create a list of unencrypted EBS volumes. Create a Systems Manager Automation runbook that includes the steps to create a new encrypted EBS volume. Modify the AWS account setting for EBS encryption to always encrypt new EBS volumes.

D. Configure the AWS Config managed rule that identifies unencrypted EBS volumes. Configure an automatic remediation action. Associate an AWS Systems Manager Automation runbook that includes the steps to create a new encrypted EBS volume. Modify the AWS account setting for EBS encryption to always encrypt new EBS volumes.

ANSWER: D

Explanation:

Configure the AWS Config managed rule that identifies unencrypted EBS volumes. Configure an automatic remediation action. Associate an AWS Systems Manager Automation runbook that includes the steps to create a new encrypted EBS volume. Modify the AWS account setting for EBS encryption to always encrypt new EBS volumes.

The AWS Config managed rule `encrypted-volumes` evaluates attached Amazon EBS volumes and reports those that are not encrypted as noncompliant. AWS Config remediation actions can invoke an AWS Systems Manager Automation

runbook, providing a scalable, event-driven mechanism to remediate the existing noncompliant fleet. The automation workflow can create a snapshot of an affected volume, create an encrypted volume from that snapshot, and coordinate attachment of the replacement volume to the EC2 instance. Operational workflows involving root volumes may require controlled instance stop/start steps, which can also be handled by the runbook.

Enabling EBS encryption by default is an account- and Region-level setting that automatically encrypts newly created EBS volumes and copied snapshots. This establishes the preventative control required for future development activity, without relying on each team to specify encryption settings for every volume request. Together, continuous AWS Config compliance evaluation and automatic remediation address existing volumes, while encryption by default prevents future unencrypted volume creation. See AWS Config encrypted-volumes rule documentation and Amazon EBS encryption by default documentation.

QUESTION NO: 55

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 ¡¢¡' 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 VPC. Attach 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 image on the EFS volume.

ANSWER: C

Explanation:

Set up an S3 gateway VPC endpoint in the VPC. Attach an endpoint policy to the endpoint to allow the required actions on the S3 bucket. This allows the EC2 instances in private subnets to access Amazon S3 directly through the AWS network rather than routing S3 requests through the NAT gateways. NAT gateways incur data-processing charges, so avoiding them for the service's substantial daily S3 retrieval traffic can materially reduce ongoing costs.

An S3 gateway endpoint is a highly available, horizontally scaled VPC component that AWS manages. It has no hourly charge and does not require the operational work associated with maintaining, patching, scaling, or replacing network appliances. The endpoint is associated with the private-subnet route tables, enabling S3 traffic to use the endpoint while preserving the existing private-subnet design.

The endpoint policy provides an additional resource-level control point, allowing access to be restricted to only the required S3 bucket and actions. This supports the stated security posture by keeping the workloads private and ensuring that S3 access is explicitly scoped. The S3 bucket policy can also use the `aws:SourceVpce` condition key to require requests through the designated endpoint. See the AWS documentation on Amazon S3 gateway endpoints and restricting S3 bucket access to a VPC endpoint.

QUESTION NO: 56

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.

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.

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:

Creating a new S3 bucket in us-east-1 and configuring S3 Cross-Region Replication from the eu-west-1 bucket provides a regional copy of the frequently updated weather-map objects. This is important because CloudFront can serve cached objects quickly at edge locations, but cache misses, expirations, and revalidations still require CloudFront to retrieve an object from an origin. A nearby us-east-1 origin reduces the latency of those origin fetches for North American viewers.

Using Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1 enables geographic origin selection while retaining the existing CloudFront distribution and the eu-west-1 bucket for other viewers. An origin-request Lambda@Edge function can inspect CloudFront-provided viewer-location headers and dynamically update the request origin to the appropriate regional S3 bucket. Together, replication keeps the regional origin populated and Lambda@Edge directs applicable requests to that closer origin, improving performance when CloudFront cannot satisfy a request from its cache. S3 replication requires versioning on both buckets and the appropriate replication configuration and IAM permissions. See Amazon S3 Replication and Lambda@Edge for CloudFront.

QUESTION NO: 57

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

Which solution will meet this requirement?

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

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

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

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

ANSWER: C

Explanation:

“From Example Corp’s management AWS account, send an invitation to AnyCompany’s AWS account. Accept the invitation in AnyCompany’s account. Create an OrganizationAccountAccessRole IAM role with the appropriate managed IAM policy.” is correct because an existing standalone AWS account joins an AWS organization through an invitation initiated by that organization’s management account. An administrator in the invited account must accept the invitation to complete the membership process. This is the standard AWS Organizations workflow for bringing an existing account under centralized governance, including organization policies, consolidated billing, and service integrations available to member accounts.

After acceptance, the account becomes a member account. AWS automatically creates `OrganizationAccountAccessRole` when an account is created directly through AWS Organizations, but it does not automatically create this role for an account that joins by invitation. Therefore, an administrator in AnyCompany’s account should manually create the role, configure its trust policy to allow the organization’s management account to assume it, and grant the permissions required for cross-account administration, commonly administrative permissions where appropriate. This preserves controlled administrative access while allowing the acquired account to be integrated into Example Corp’s organization. See the AWS guidance for inviting AWS accounts to join an organization and accessing member accounts.

QUESTION NO: 58

A company runs an loT platform on AWS loT sensors in various locations send data to the company ' s Node js API servers on Amazon EC2 instances running behind an Application Load Balancer The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume

The number of sensors the company has deployed in the field has increased over time and is expected to grow significantly The API servers are consistently overloaded and RDS metrics show high write latency

Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? {Select TWO.)

A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume ' s IOPS

B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance andadd read replicas

C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data

D. Use AWS X-Ray to analyze and debug application issues and add more API servers to match the load

E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

ANSWER: C E

Explanation:

“Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data” provides a durable, scalable ingestion layer that decouples high-volume sensor traffic from the API and database tiers. Kinesis Data Streams can accept records from many concurrent producers and retain them for processing, while Lambda can consume records from the stream and scale processing in response to the workload. This avoids requiring the company to continually size and operate an EC2 API fleet solely to handle fluctuating sensor-write traffic. Kinesis also provides buffering, so short-term surges do not immediately overload downstream storage. Amazon documents this streaming model at Amazon Kinesis Data Streams: What is Kinesis Data Streams?.

“Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance” supplies the write scalability needed for rapidly growing, device-generated data. DynamoDB is a fully managed NoSQL service designed for consistent low-latency reads and writes at scale. A well-designed partition key, such as a sufficiently distributed device identifier combined with an appropriate time component, spreads writes across partitions. On-demand capacity can automatically accommodate unpredictable growth without capacity planning, or provisioned capacity with auto scaling can be used when demand is more predictable. This removes the relational database storage-write bottleneck and aligns the persistence layer with the streaming architecture. See DynamoDB read/write capacity modes.

QUESTION NO: 59

A company runs a proprietary stateless ETL application on an Amazon EC2 Linux instance. The application is a Linux binary, and the source code cannot be modified. The application is single-threaded, uses 2 GB of RAM. and is highly CPU intensive The application is scheduled to run every 4 hours and runs for up to 20 minutes A solutions architect wants to revise the architecture for the solution.

Which strategy should the solutions architect use?

A. Use AWS Lambda to run the application. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours

B. Use AWS Batch to run the application Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours

C. Use AWS Fargate to run the application Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours

D. Use Amazon 6C2 Spot Instances to run the application Use AWS CodeDeptoy to deploy and run the application every 4 hours.

ANSWER: C

Explanation:

Use AWS Fargate to run the application Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours is the appropriate strategy because this is a stateless, short-duration, scheduled workload. The existing proprietary Linux binary can be placed in a container image without modifying its source code. An Amazon ECS task running on AWS Fargate then provides the required Linux execution environment while eliminating the need to maintain an always-on EC2 instance, including host provisioning, patching, and capacity management.

A task definition can allocate a supported Fargate CPU and memory combination that includes 2 GB of memory and sufficient CPU capacity for the CPU-intensive, single-threaded process. Since the task runs only for up to 20 minutes every four hours, Fargate allows compute capacity to exist only during execution, which is operationally simple and avoids paying for an idle instance between runs. Amazon EventBridge supports scheduled rules that invoke Amazon ECS tasks at recurring intervals, making it suitable for launching the task every four hours. See the AWS Fargate documentation and Amazon ECS scheduled tasks with EventBridge Scheduler documentation.

QUESTION NO: 60

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-tier 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? (Choose 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:

A Route 53 private hosted zone answers DNS queries only for VPCs that are associated with that hosted zone. Because the application runs in a new VPC owned by Account B while the private hosted zone is owned by Account A, the VPC must be

explicitly associated across accounts before the EC2 instance can resolve `db.example.com`. The hosted-zone-owning account first creates an authorization for the specific Account B VPC. This authorization grants the VPC owner permission to complete the cross-account association without transferring ownership of the hosted zone. Account B then associates its new VPC with the hosted zone in Account A. Once the association is established, resources using the VPC's Amazon-provided DNS can resolve records in that private hosted zone, including the CNAME that points to the RDS endpoint. The association authorization is only required to establish the cross-account association, so it can be deleted afterward as cleanup. This preserves centralized management of private DNS records in Account A while enabling private name resolution for the workload VPC in Account B. See the AWS guidance for associating a private hosted zone with VPCs in different AWS accounts and the CreateVPCAssociationAuthorization API reference.

QUESTION NO: 61

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 AWSLambda 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:

"Deploy the application front end to an Amazon S3 bucket served by Amazon CloudFront Deploy the application backend using Amazon API Gateway with an AWSLambda proxy integration" provides a serverless, highly available architecture with minimal infrastructure management. Amazon S3 and CloudFront can host and deliver the mobile application's static front-end content at scale, while API Gateway automatically handles request routing, throttling, and scaling for the backend API. Lambda scales in response to the incoming Friday submission volume without requiring capacity planning, servers, load balancers, or scheduled scaling actions.

"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." supplies a low-operational-overhead reporting solution. S3 provides durable and highly available object storage for the submitted records. Data can be stored in analytics-friendly formats and queried directly with Athena, a serverless query service that scales according to query demand. QuickSight can use Athena as its data source to create and publish monthly payroll reports without maintaining a database cluster or reporting servers. Together, these services satisfy the variable ingestion and reporting workloads through managed, elastic AWS services. See Amazon API Gateway Lambda integrations and What is Amazon Athena.

QUESTION NO: 62

A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.

The company has the following DNS resolution requirements:

• On-premises systems should be able to resolve and connect to cloud.example.com.

• All VPCs should be able to resolve cloud.example.com.

There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway. Which architecture should the company use to meet these requirements with the HIGHEST performance?

A. Associate the private hosted zone to all the VPCs. Create a Route 53 inbound resolver in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

B. Associate the private hosted zone to all the VPCs. Deploy an Amazon EC2 conditional forwarder in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the conditional forwarder.

C. Associate the private hosted zone to the shared services VPC. Create a Route 53 outbound resolver in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the outbound resolver.

D. Associate the private hosted zone to the shared services VPC. Create a Route 53 inbound resolver in the shared services VPC. Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

ANSWER: A

Explanation:

Associate the private hosted zone to all the VPCs. Create a Route 53 inbound resolver in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver. This design provides managed, highly available hybrid DNS resolution without introducing self-managed DNS forwarder instances. Associating the Route 53 private hosted zone with every applicable VPC allows Amazon-provided DNS resolution within each VPC to return the private records for `cloud.example.com`. A Route 53 Resolver inbound endpoint supplies IP addresses in the shared services VPC that on-premises DNS servers can use as conditional-forwarding targets. Queries for the private domain therefore enter AWS through the inbound endpoint and are resolved by Route 53 Resolver against the private hosted zone. The existing Direct Connect and Transit Gateway architecture provides the private network path for both DNS traffic to the endpoint and application connectivity from on premises to resources in the attached VPCs. Route 53 Resolver endpoints are an AWS managed service and are deployed using endpoint IP addresses in multiple subnets for resilience; AWS recommends configuring on-premises DNS forwarding to those inbound endpoint IPs. This approach meets the cross-network name-resolution requirement while retaining native private hosted zone behavior across all VPCs. See Route 53 Resolver inbound endpoints and associating VPCs with a private hosted zone.

QUESTION NO: 63

A research center is migrating to the AWS Cloud and has moved its on-premises 1 PB object storage to an Amazon S3 bucket. One hundred scientists are using this object storage to store their work-related documents. Each scientist has a personal folder on the object store. All the scientists are members of a single IAM user group. The research center's compliance officer is worried that scientists will be able to access each other's work. The research center has a strict obligation to report on which scientist accesses which documents. The team that is responsible for these reports has little AWS experience and wants a ready-to-use solution that minimizes operational overhead. Which combination of actions should a solutions architect take to meet these requirements? (Choose two.)

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

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

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

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

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

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

ANSWER: B F

Explanation:

Creating an identity policy that grants the user read and write access, with S3 paths prefixed by `${aws:username}`, provides scalable per-scientist isolation while allowing the same policy to be attached to the shared IAM group. IAM policy variables are evaluated at request time using the authenticated IAM user's attributes. Therefore, a resource pattern such as `arn:aws:s3:::bucket-name/${aws:username}/*` dynamically resolves to each scientist's own folder. The policy should also allow the required bucket-level actions, such as listing the bucket with an appropriate prefix condition, so users can work within their assigned prefix. AWS documents this dynamic authorization approach in its IAM policy variables documentation.

Configuring an AWS CloudTrail trail to capture S3 object-level data events supplies the audit record required to determine which authenticated principal accessed specific S3 objects and when. CloudTrail data events include S3 object API activity, such as object reads and writes, that is not included in standard management-event logging. Delivering the trail to a separate S3 bucket creates durable, centrally stored log files. Amazon Athena can then query those files directly using SQL, enabling the reporting team to generate access reports without operating custom log-collection infrastructure. See Logging Amazon S3 API calls using AWS CloudTrail.

QUESTION NO: 64

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:

Creating an S3 Lifecycle configuration to expire incomplete multipart uploads 7 days after initiation reduces storage charges from failed or abandoned uploads. With multipart upload, Amazon S3 stores each successfully uploaded part until the multipart upload is completed or explicitly aborted. Poor connectivity makes incomplete uploads likely, so automatically aborting them after an appropriate retention period removes unused parts and stops their continuing storage charges. Amazon S3 Lifecycle supports an AbortIncompleteMultipartUpload action specifically for this purpose. AWS documents this behavior at Aborting incomplete multipart uploads using a bucket lifecycle configuration.

Creating an S3 Lifecycle configuration to transition objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 180 days aligns the storage class with the stated access pattern. S3 Standard provides the appropriate performance and availability characteristics while videos are frequently viewed during their first 180 days. After that period, S3 Standard-IA provides the same millisecond retrieval performance for occasional access at a lower storage price, with retrieval and minimum-duration pricing that are appropriate for rarely accessed video objects. The videos are substantially larger than the

128 KB minimum billable object size for S3 Standard-IA, supporting the expected storage-cost benefit. See Amazon S3 Lifecycle transition considerations.

QUESTION NO: 65

A company is storing sensitive data in an Amazon S3 bucket. The company must log all activities for objects in the S3 bucket and must keep the logs for 5 years. The company's security team also must receive an email notification every time there is an attempt to delete data in the S3 bucket.

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

A. Configure AWS CloudTrail to log S3 data events.

B. Configure S3 server access logging for the S3 bucket.

C. Configure Amazon S3 to send object deletion events to Amazon Simple Email Service (Amazon SES).

D. Configure an Amazon EventBridge rule for CloudTrail DeleteObject data events 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 because S3 object-level operations, including `GetObject`, `PutObject`, and `DeleteObject`, are CloudTrail data events. This provides an audit record of activity involving the objects in the protected bucket, including delete API calls and their outcomes. Configure an Amazon EventBridge rule that matches CloudTrail `DeleteObject` API-call events and publishes to an Amazon SNS topic. An SNS topic can deliver notifications to an email subscription, allowing the security team to receive an alert for each attempted deletion, including an API call that does not succeed. Configure a new S3 bucket to store the CloudTrail log files with an S3 Lifecycle policy. The lifecycle configuration can retain the logs for the required five-year period while transitioning older log files to lower-cost archival S3 storage classes, such as S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive, as access needs allow. This separates audit logs from the source data and provides durable, low-cost retention. See CloudTrail data events documentation and Amazon S3 Lifecycle documentation.

QUESTION NO: 66

A company runs a Java application that has complex dependencies on VMs that are in the company ' s data center. The application is stable. but the company wants to modernize the technology stack. The company wants to migrate the application to AWS and minimize the administrative overhead to maintain the servers.

Which solution will meet these requirements with the LEAST code changes?

A. Migrate the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate by using AWS App2Container. Store container images in Amazon Elastic Container Registry (Amazon ECR). Grant the ECS task execution role permission to access the ECR image repository. Configure Amazon ECS to use an Application Load Balancer (ALB). Use the ALB to interact with the application.

B. Migrate the application code to a container that runs in AWS Lambda. Build an Amazon API Gateway REST API with Lambda integration. Use API Gateway to interact with the application.

C. Migrate the application to Amazon Elastic Kubernetes Service (Amazon EKS) on EKS managed node groups by using AWS App2Container. Store container images in Amazon Elastic Container Registry (Amazon ECR). Give the EKS nodes permission to access the ECR image repository. Use Amazon API Gateway to interact with the application.

D. Migrate the application code to a container that runs in AWS Lambda. Configure Lambda to use an Application Load Balancer (ALB). Use the ALB to interact with the application.

ANSWER: A

Explanation:

Migrate the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate by using AWS App2Container because it combines a low-code-change modernization path with serverless container operations. AWS App2Container is designed to analyze existing Java applications running on virtual machines or physical servers, identify application artifacts and dependencies, and containerize the application. This allows the company to retain the stable application's existing runtime model rather than redesigning it around a new event-driven or function-based architecture. The resulting image can be stored in Amazon ECR, and the ECS task execution role provides the task with the permissions required to retrieve that image when it starts.

Amazon ECS with the Fargate launch type is especially appropriate for the operational requirement. Fargate runs containers without the company having to provision, patch, scale, or manage the underlying EC2 worker instances. ECS manages task placement and integrates directly with an Application Load Balancer, which can distribute HTTP or HTTPS requests to healthy application tasks and provide a stable entry point for clients. Consequently, this approach modernizes delivery and deployment through containers while preserving the application with minimal code changes and minimizing ongoing server administration. See the AWS App2Container User Guide and the Amazon ECS on AWS Fargate documentation.

QUESTION NO: 67

A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.

The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.

Which storage strategy is the MOST cost-effective and meets the design requirements?

A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieval. Configure a lifecycle policy to delete data older than 120 days.

B. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scale. Configure the DynamoDB Time to Live (TTL) feature to delete records older than 120 days.

C. Design the application to store each incoming record in a single table in an Amazon RDS MySQL database. Run a nightly cron job that executes a query to delete any records older than 120 days.

D. Design the application to batch incoming records before writing them to an Amazon S3 bucket. Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the data. Configure a lifecycle policy to delete the data after 120 days.

ANSWER: B

Explanation:

Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scale. Configure the DynamoDB Time to Live (TTL) feature to delete records older than 120 days. This approach directly aligns with the workload's combination of very high write volume, sub-4 KB records, durable storage, and consistently low-latency retrieval. DynamoDB is a fully managed NoSQL database built for key-value access at virtually any scale, and it replicates data across multiple Availability Zones in a Region for durability. A well-designed, high-cardinality partition key distributes writes across partitions, allowing the table to accommodate ingestion from a large global device fleet without concentrating traffic on a small number of keys.

DynamoDB capacity calculations use item size, so these small records are an efficient fit for the service. The application can select on-demand capacity for unpredictable traffic or provisioned capacity with auto scaling when throughput is known, enabling cost control while meeting demand. TTL is the appropriate retention mechanism: each item receives an expiration-time attribute, and DynamoDB automatically removes expired items without application-managed deletion jobs. This continuously limits retained data to the required 120-day window and avoids operational overhead. AWS documents DynamoDB's low-latency, scalable data model in the Amazon DynamoDB Developer Guide and explains automatic expiry in the DynamoDB Time to Live documentation.

QUESTION NO: 68

An entertainment company hosts a ticketing service on a fleet of Linux Amazon EC2 instances that are in an Auto Scaling group. The ticketing service uses a pricing file. The pricing file is stored in an Amazon S3 bucket that has S3 Standard storage. A central pricing solution that is hosted by a third party updates the pricing file. The pricing file is updated every 1-15 minutes and has several thousand line items. The pricing file is downloaded to each EC2 instance when the instance launches. The EC2 instances occasionally use outdated pricing information that can result in incorrect charges for customers. Which solution will resolve this problem MOST cost-effectively?

A. Create an AWS Lambda function to update an Amazon DynamoDB table with new prices each time the pricing file is updated. Update the ticketing service to use DynramoDB to look up pricing

B. Create an AWS Lambda function to update an Amazon Elastic File System (Amazon EFS) file share with the pricing file each time the file is updated. Update the ticketing service to use Amazon EFS to access the pricing file.

C. Load Mountpoint for Amazon S3 onto the AMI of the EC2 instances. Configure Mountpoint for Amazon S3 to mount the S3 bucket that contains the pricing file. Update the ticketing service to point to the mount point and path to access the S3 object.

D. Create an Amazon Elastic Block Store (Amazon EBS) volume. Use EBS Multi-Attach to attach the volume to every EC2 instance. When a new EC2 instance launches, configure the new instance to update the pricing file on the EBS volume. Update the ticketing service to point to the new local source.

ANSWER: C

Explanation:

Loading Mountpoint for Amazon S3 onto the EC2 AMI, mounting the bucket, and changing the ticketing service to read the pricing file from the mount path is the most cost-effective solution. It eliminates the one-time, instance-launch download that causes stale local copies. Instead, every instance accesses the centrally maintained S3 object through a file-system-style path, which allows the existing Linux-based service to use the current shared source without requiring a separate distributed file system or a pricing-data transformation pipeline.

Mountpoint for Amazon S3 is an AWS-supported, high-throughput file client for Linux that presents S3 objects through a local mount point. This is well suited to a centrally updated pricing file that instances primarily read. Amazon S3 also provides strong read-after-write consistency: after a successful overwrite or PUT, subsequent reads return the latest object version. In practice, the application should open or reread the file when it needs refreshed pricing, and any Mountpoint caching configuration should be selected to meet the required freshness interval. The solution retains S3 Standard as the durable central store and adds only lightweight client configuration to the fleet. See Mountpoint for Amazon S3 and Amazon S3 data consistency model.

QUESTION NO: 69

A company has an organization in AWS Organizations. The company is using AWS Control Tower to deploy a landing zone for the organization. The company wants to implement governance and policy enforcement. The company must implement a policy that will detect Amazon RDS DB instances that are not encrypted at rest in the company's production OU.

Which solution will meet this requirement?

A. Turn on mandatory guardrails in AWS Control Tower. Apply the mandatory guardrails to the production OU.

B. Enable the appropriate guardrail from the list of strongly recommended guardrails in AWS Control Tower. Apply the guardrail to the production OU.

C. Use AWS Config to create a new mandatory guardrail. Apply the rule to all accounts in the production OU.

D. Create a custom SCP in AWS Control Tower. Apply the SCP to the production OU.

ANSWER: B

Explanation:

Enable the appropriate guardrail from the list of strongly recommended guardrails in AWS Control Tower. Apply the guardrail to the production OU. AWS Control Tower provides managed controls (formerly called guardrails) that can be enabled at the organizational-unit level. For this requirement, the relevant capability is a *detective* control that continuously evaluates whether Amazon RDS DB instances use encryption at rest. Detective controls use AWS Config rules to assess deployed resources and report compliance status when resources do not meet the configured governance requirement.

Applying this control to the production OU ensures that it is deployed consistently to the enrolled member accounts governed by that OU. This provides centralized visibility of noncompliant RDS instances through AWS Control Tower and AWS Config, while preserving the requested detection-focused approach. Control Tower manages the lifecycle and OU-level inheritance of the enabled control, avoiding the need to independently deploy and maintain equivalent Config rules in every production account. AWS documents the available Control Tower controls and their behavior in the AWS Control Tower controls reference. The process and scope for applying a control to an OU are described in Enable controls in AWS Control Tower.

QUESTION NO: 70

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? (Select THRE

E.

)

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.

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.

F. Implement least privilege access for the IAM service role that is assigned to AWS Backup.

ANSWER: A B C

Explanation:

“Implement cross-account backup with AWS Backup vaults in designated non-production accounts” creates an administrative boundary between production workloads and their recovery copies. A compromised privileged identity in a production account does not automatically gain authority over backup vaults and recovery points held in a separate, dedicated backup account. AWS Backup supports cross-account backup when accounts belong to the same AWS Organization, making this approach suitable for a large multi-account environment.

“Add an SCP that restricts the modification of AWS Backup vaults” provides organization-level preventive guardrails. Service control policies set the maximum permissions available in member accounts, including for highly privileged identities, so the organization can deny actions that would delete or alter backup protections. This helps ensure that production-account credentials cannot be used to weaken the backup controls.

“Implement AWS Backup Vault Lock in compliance mode” makes recovery points immutable for their configured retention periods. After the Vault Lock grace period expires, even users with administrative permissions cannot delete or alter protected recovery points before retention expires. Together, account isolation, organization-level restrictions, and immutable retention protect backups against ransomware operators who obtain privileged credentials. See AWS Backup cross-account backup documentation and AWS Backup Vault Lock documentation.

QUESTION NO: 71

A company wants to migrate its on-premises data center to the AWS Cloud. This includes thousands of virtualized Linux and Microsoft Windows servers, SAN storage, Java and PHP applications with MySQL, and Oracle databases. There are many dependent services hosted either in the same data center or externally. The technical documentation is incomplete and outdated. A solutions architect needs to understand the current environment and estimate the cloud resource costs after the migration. Which tools or services should the solutions architect use to plan the cloud migration? (Choose three.)

A. AWS Application Discovery Service

B. AWS SMS

C. AWS X-Ray

D. AWS Cloud Adoption Readiness Tool (CART)

E. Amazon Inspector

F. AWS Migration Hub

ANSWER: A D F

Explanation:

AWS Application Discovery Service is appropriate because it gathers information from the on-premises estate when existing documentation is incomplete. Its agents and agentless collector can capture server inventory, system configuration, utilization metrics, and network connection data. This discovery data establishes an evidence-based view of the server fleet, application dependencies, and resource consumption needed for migration wave planning and right-sizing AWS target resources.

AWS Migration Hub provides the central migration-planning workspace for discovered servers and applications. It integrates discovery information, enables application grouping, and helps track and coordinate migrations across AWS migration services and partner tools. This supports organizing a large, interdependent portfolio into manageable migration activities.

AWS Cloud Adoption Readiness Tool (CART) complements technical discovery by assessing organizational cloud-adoption readiness. It identifies readiness gaps across dimensions such as business, people, governance, platform, security, and operations, enabling the company to include necessary organizational remediation in its migration plan. Together, these services address estate discovery, migration coordination, and cloud-readiness planning. See AWS Application Discovery Service documentation and AWS Migration Hub documentation.

QUESTION NO: 72

A media storage application uploads user photos to Amazon S3 for processing by AWS Lambda functions. Application state is stored in Amazon DynamoDB tables. Users are reporting that some uploaded photos are not being processed properly. The application developers trace the logs and find that Lambda is experiencing photo processing issues when thousands of users upload photos simultaneously. The issues are the result of Lambda concurrency limits and the performance of DynamoDB when data is saved. Which combination of actions should a solutions architect take to increase the performance and reliability of the application? (Choose two.)

A. Evaluate and adjust the RCUs for the DynamoDB tables.

B. Evaluate and adjust the WCUs for the DynamoDB tables.

C. Add an Amazon ElastiCache layer to increase the performance of Lambda functions.

D. Add an Amazon Simple Queue Service (Amazon SQS) queue and reprocessing logic between Amazon S3 and the Lambda functions.

E. Use S3 Transfer Acceleration to provide lower latency to users.

ANSWER: B D

Explanation:

“Evaluate and adjust the WCUs for the DynamoDB tables” addresses the write workload created when many concurrent photo-processing invocations save application state. In provisioned capacity mode, DynamoDB write capacity units define the write throughput available to a table. Sizing write capacity appropriately, and configuring DynamoDB auto scaling where applicable, allows the table to accommodate increased concurrent writes while avoiding throttling that can interrupt or delay processing. DynamoDB capacity-mode and throughput guidance is available in the DynamoDB Developer Guide.

“Add an Amazon Simple Queue Service (Amazon SQS) queue and reprocessing logic between Amazon S3 and the Lambda functions” decouples object uploads from photo processing. Amazon S3 event notifications can deliver event messages to an SQS queue, and Lambda polls that queue in batches. The queue durably buffers a sudden upload burst, while Lambda scales processing at a controlled rate rather than requiring every upload to be handled immediately. This pattern improves resilience through Lambda and SQS retry behavior, and a dead-letter queue or redrive policy can retain messages requiring reprocessing. AWS documents the polling, batching, scaling, and error-handling behavior in Using Lambda with Amazon SQS.

QUESTION NO: 73

A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.

The solutions architect created the following IAM policy and attached it to an IAM role:

During tests, me solutions architect was able to successfully get existing test objects m the S3 bucket However, attempts to upload a new object resulted in an error message. The error message stated that me action was forbidden.

Which action must me solutions architect add to the IAM policy to meet all the requirements?

A. Kms:GenerateDataKey

B. KmsGetKeyPolpcy

C. kmsGetPubKKey

D. kms:SKjn

ANSWER: A

Explanation:

`kms:GenerateDataKey` is required because client-side encryption using an AWS KMS customer managed key uses envelope encryption. Before uploading a new object, the client must request a unique data key from AWS KMS under the selected KMS key. AWS KMS returns both a plaintext data key and an encrypted copy of that data key. The encryption client uses the plaintext key locally to encrypt the object content before sending the encrypted object to Amazon S3, while it stores the encrypted data key with the object's encryption metadata.

The role's ability to retrieve existing encrypted objects indicates it has the permissions needed to obtain and decrypt the encrypted data key associated with those objects. A new encrypted upload, however, requires AWS KMS to generate a fresh data key. Adding `kms:GenerateDataKey` for the applicable KMS key authorizes that operation without exposing the KMS key material itself. AWS KMS keeps the customer managed key inside the service and uses it to protect the generated data key. This is the standard permission needed for applications that encrypt new data locally with KMS-backed envelope encryption. See the AWS KMS data keys documentation and the Amazon S3 client-side encryption guide.

QUESTION NO: 74

A company runs a file-processing application in one AWS Region. The application stores shared files in Amazon Elastic File System (Amazon EFS). The company needs a disaster recovery solution in a second Region. The company requires an RPO of less than 1 hour and wants to minimize operational overhead. A recovery environment already exists in the second Region.

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

A. Configure Amazon EFS replication from the source file system to an EFS file system in the recovery Region.

B. Configure Amazon Route 53 failover routing records for the primary and recovery application endpoints. Associate a health check with the primary record.

C. Create an Amazon EFS mount target in the recovery Region for the existing source file system.

D. Use Amazon S3 Cross-Region Replication to replicate the EFS mount targets to the recovery Region.

E. Create an Amazon EBS snapshot of the EFS file system every hour and restore the snapshot during a disaster.

ANSWER: A B

Explanation:

Configure Amazon EFS replication from the source file system to a file system in the recovery Region. EFS replication continuously and asynchronously replicates data to a read-only replica file system in another Region. The service is managed and designed to provide a low RPO without requiring the company to build file synchronization software or operate replication servers.

Configure Amazon Route 53 failover routing records for the application endpoints in the primary and recovery Regions. Associate a health check with the primary record. During a failure, Route 53 can return the recovery endpoint for new DNS resolutions. The recovery procedure must promote the EFS replica so that it becomes writable before the recovery application processes new files. This combination provides managed data replication and DNS-based traffic redirection. See Amazon EFS replication and Amazon Route 53 failover routing.

QUESTION NO: 75

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 visualization 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. The AWS Application Discovery Agent is designed for detailed on-premises server discovery during migration planning. After installation on each virtual machine, it collects the information needed to right-size target Amazon EC2 instances, including system configuration and time-series utilization data for CPU, memory, and disk activity. It also provides an inventory of running processes and captures network connection information, allowing the company to identify server dependencies and map communication patterns before migration.

This agent-based approach directly satisfies both the performance-sizing and application-dependency requirements from a single AWS migration-discovery service. The collected data can be viewed and analyzed through AWS Application Discovery Service and can support migration planning workflows, including use with AWS Migration Hub. AWS documents that the agent gathers detailed system-performance, process, and network-connection data from on-premises servers. See the AWS Application Discovery Agent documentation and the AWS Application Discovery Service User Guide.

QUESTION NO: 76

A company uses AWS Organizations to manage hundreds of AWS accounts. The company must collect AWS API activity from all existing and future accounts into a centralized log archive account. The security team requires that CloudTrail log files cannot be deleted or changed for 7 years, including by administrators in member accounts.

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

A. Create an organization trail from the AWS Organizations management account. Deliver the trail logs to an Amazon S3 bucket in the log archive account.

B. Enable S3 Object Lock in compliance mode on the log archive bucket. Configure a default retention period of 7 years.

C. Create and attach an SCP to member account OUs that denies cloudtrail:StopLogging, cloudtrail:DeleteTrail, and cloudtrail:UpdateTrail.

D. Create an AWS CloudTrail trail separately in every member account. Deliver the logs to an S3 bucket in the same member account.

E. Enable S3 versioning on the log archive bucket without enabling S3 Object Lock.

F. Create Amazon CloudWatch alarms for CloudTrail API calls and invoke an AWS Lambda function to recreate deleted log files.

ANSWER: A B C

Explanation:

An organization trail that is created from the AWS Organizations management account records events for all accounts in the organization, including accounts that are added later. Configuring the trail to deliver logs to an Amazon S3 bucket in a dedicated log archive account centralizes the audit data outside the member accounts that generate it.

Amazon S3 Object Lock in compliance mode prevents protected object versions from being overwritten or deleted until their retention period expires. Configuring a default retention period of 7 years on the destination bucket provides immutable storage for the CloudTrail log files. An SCP that denies CloudTrail modification and stop-logging actions in member accounts

prevents member-account administrators from disabling or changing the organization-wide logging configuration. For more information, see Creating an organization trail and Using S3 Object Lock.

QUESTION NO: 77

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, provides Availability Zone resilience for the inspection tier. An appliance failure or an Availability Zone impairment does not remove all firewall capacity, which supports highly available centralized egress inspection within the Region.

Creating a Gateway Load Balancer with a target group containing the firewall appliance instances is the purpose-built AWS pattern for transparently inserting and scaling third-party virtual network appliances, including firewalls. Gateway Load Balancer distributes traffic to healthy registered appliances and uses flow-based processing that preserves traffic symmetry, a key requirement for stateful firewall sessions. Its health checks enable unhealthy appliances to be removed from service without requiring route-table updates or a manual failover process.

Creating a VPC Gateway Load Balancer endpoint and making it the next hop in the shared services VPC route table steers traffic to the Gateway Load Balancer for inspection. Gateway Load Balancer endpoints are route-table targets specifically designed for this appliance-insertion architecture. Together, multi-AZ appliances, Gateway Load Balancer health-based distribution, and endpoint routing provide centralized inspection with fast automated failover. See the AWS Gateway Load Balancer documentation and Gateway Load Balancer endpoint documentation.

QUESTION NO: 78

A retail company is operating its ecommerce application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses an Amazon RDS DB instance as the database backend. Amazon CloudFront is configured with one origin that points to the ALB. Static content is cached. Amazon Route 53 is used to host all public zones. After an update of the application, the ALB occasionally returns a 502 status code (Bad Gateway) error. The root cause is malformed HTTP headers that are returned to the ALB. The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs. While the company is working on the problem, the solutions architect needs to provide a custom error page instead of the standard ALB error page to visitors. Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Choose two.)

A. Create an Amazon S3 bucket. Configure the S3 bucket to host a static webpage. Upload the custom error pages to Amazon S3.

B. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Target.FailedHealthChecks is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a publicly accessible web server.

C. Modify the existing Amazon Route 53 records by adding health checks. Configure a fallback target if the health check fails. Modify DNS records to point to a publicly accessible webpage.

D. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Elb.InternalError is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a public accessible web server.

E. Add an Amazon S3 origin to the CloudFront distribution and configure a CloudFront custom error response to return the S3-hosted error page for HTTP 502 responses.

ANSWER: A E

Explanation:

The appropriate low-overhead solution is to store a static, branded error document in Amazon S3 and have CloudFront serve that document when the application origin returns HTTP 502. **Create an Amazon S3 bucket. Configure the S3 bucket to host a static webpage. Upload the custom error pages to Amazon S3.** provides a durable managed location for the error asset without requiring servers, application deployments, health-check automation, or DNS failover. The bucket should be available to CloudFront through an S3 origin, using appropriate access controls.

Add an Amazon S3 origin to the CloudFront distribution and configure a CloudFront custom error response to return the S3-hosted error page for HTTP 502 responses. centralizes error handling at the layer that already receives viewer requests. CloudFront can map a 502 response from the ALB origin to a configured response page path and optionally specify the response code and error-caching duration. This means visitors receive the custom page during intermittent ALB-origin failures, while normal requests continue to use the ALB origin after recovery. AWS documents the supported error codes, response-page paths, and caching behavior in the CloudFront custom error responses documentation. For static error content, see Hosting a static website using Amazon S3.

QUESTION NO: 79

A company has five development teams that have each created five AWS accounts to develop and host applications. To track spending, the development teams log in to each account every month, record the current cost from the AWS Billing and Cost Management console, and provide the information to the company ' s finance team.

The company has strict compliance requirements and needs to ensure that resources are created only in AWS Regions in the United States. However, some resources have been created in other Regions.

A solutions architect needs to implement a solution that gives the finance team the ability to track and consolidate expenditures for all the accounts. The solution also must ensure that the company can create resources only in Regions in the United States.

Which combination of steps will meet these requirements in the MOST operationally efficient way? (Select THREE.)

A. Create a new account to serve as a management account. Create an Amazon S3 bucket for the finance team. Use AWS Cost and Usage Reports to create monthly reports and to store the data in the finance team's S3 bucket.

B. Create a new account to serve as a management account. Deploy an organization in AWS Organizations with all features enabled. Invite all the existing accounts to the organization. Ensure that each account accepts the invitation.

C. Create an OU that includes all the development teams. Create an SCP that allows the creation of resources only in Regions that are in the United States. Apply the SCP to the OU.

D. Create an OU that includes all the development teams. Create an SCP that denies the creation of resources in Regions that are outside the United States. Apply the SCP to the OU.

E. Create an IAM role in the management account. Attach a policy that includes permissions to view the Billing and Cost Management console. Allow the finance team users to assume the role. Use AWS Cost Explorer and the Billing and Cost Management console to analyze costs.

F. Create an IAM role in each AWS account. Attach a policy that includes permissions to view the Billing and Cost Management console. Allow the finance team users to assume the role.

ANSWER: B D E

Explanation:

Creating a new account to serve as a management account, enabling all AWS Organizations features, and inviting the existing accounts establishes centralized multi-account governance and consolidated billing. The management account can view charges incurred by all member accounts, eliminating the need for teams to collect costs manually from each account. AWS Organizations also provides the policy framework required to apply preventive controls consistently across the company's accounts.

Creating an OU that includes all development accounts and applying an SCP that denies resource-creation requests outside United States Regions enforces the required geographic restriction centrally. A deny-based SCP using the `aws:RequestedRegion` condition establishes a permission boundary that member-account IAM policies cannot override. The SCP should account for required global services and explicitly allow the supported US Regions.

Allowing finance users to assume an IAM role in the management account with Billing and Cost Management permissions gives the finance team one controlled access point for consolidated cost analysis. From that account, the team can use AWS Cost Explorer and the Billing and Cost Management console to review organization-wide spending without maintaining roles or performing monthly sign-ins across every development account. See the AWS Organizations User Guide and the AWS SCP examples.

QUESTION NO: 80

A company has automated the nightly retraining of its machine learning models by using AWS Step Functions. The workflow consists of multiple steps that use AWS Lambda Each step can fail for various reasons and any failure causes a failure of the overall workflow

A review reveals that the retraining has failed multiple nights in a row without the company noticing the failure A solutions architect needs to improve the workflow so that notifications are sent for all types of failures in the retraining process

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

A. Create an Amazon Simple Notification Service (Amazon SNS) topic with a subscription of type "Email" that targets the team's mailing list.

B. Create a task named "Email" that forwards the input arguments to the SNS topic.

C. Add a Catch field to all Task, Map, and Parallel states with "ErrorEquals": ["States.ALL"] and "Next": "Email".

D. Add a new email address to Amazon Simple Email Service (Amazon SES). Verify the email address.

E. Create a task named "Email" that forwards the input arguments to the SES email address.

F. Add a Catch field to all Task, Map, and Parallel states with "ErrorEquals": ["States.Runtime"] and "Next": "Email".

ANSWER: A B C

Explanation:

Creating an Amazon Simple Notification Service (Amazon SNS) topic with an email subscription provides a durable notification endpoint for the team's mailing list. Amazon SNS sends a confirmation message to the subscribed email address, and, after confirmation, publishes delivered to the topic are sent to the team automatically.

Creating a task named "Email" that forwards the input arguments to the SNS topic gives the Step Functions workflow a centralized notification state. This task can use the optimized Step Functions integration for Amazon SNS publishing,

allowing the state machine to publish an error payload, execution context, or other relevant failure details without requiring a separate Lambda function solely for notification delivery.

Adding a `Catch` field to every relevant Task, Map, and Parallel state with `"ErrorEquals": ["States.ALL"]` and `"Next": "Email"` routes handled state failures to that centralized notification task. Catchers let Step Functions transition to recovery or notification logic when an error occurs, and `States.ALL` is the wildcard matcher used to handle the errors that a state can catch. Together, the SNS topic, notification task, and catch handlers ensure failures in the retraining workflow generate an email notification. See the AWS Step Functions error handling documentation and the Step Functions Amazon SNS integration documentation.

QUESTION NO: 81

Question:

An application uses CloudFront, App Runner, and two S3 buckets — one for static assets and one for user-uploaded content. User content is infrequently accessed after 30 days. Users are located only in Europe.

How can the companyoptimize cost?

A. Expire S3 objects after 30 days.

B. Transition S3 content toGlacier Deep Archiveafter 30 days.

C. Use Spot Instances with App Runner.

D. Add auto scaling to Aurora read replica.

E. Use CloudFront Price Class 100 for European users.

ANSWER: B E

Explanation:

Transitioning user-uploaded S3 content to Amazon S3 Glacier Deep Archive after 30 days can substantially reduce ongoing storage charges for content that is rarely needed. An S3 Lifecycle configuration automates this transition based on object age, avoiding operational effort while retaining the objects rather than deleting them. Glacier Deep Archive is designed for long-term archival data and has the lowest S3 storage price, making it appropriate when delayed retrieval is acceptable. The design should account for its archival retrieval characteristics and minimum storage duration when setting the lifecycle policy.

Using CloudFront Price Class 100 is also appropriate because all viewers are in Europe. Price Class 100 serves Europe and lower-cost North American edge locations while excluding higher-cost edge locations elsewhere. Since the application has no users outside Europe, excluding those locations reduces CloudFront data-transfer charges without reducing performance for the stated audience. CloudFront price classes can be selected per distribution, allowing the company to align edge-location coverage with actual viewer geography. See the Amazon S3 Lifecycle transition documentation and the CloudFront price class documentation.

QUESTION NO: 82

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 monthlyrequirements.

* 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 this requirement with the LEAST amount of operational overhead? (Select 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.

G. Configure an Amazon EventBridge rule for AWS Backup Backup Job State Change events with a FAILED status, and set an Amazon SNS topic as the target.

ANSWER: A B G

Explanation:

Create an AWS Backup plan with a backup rule for each of the retention requirements. is correct because AWS Backup plans can contain multiple backup rules, each with its own schedule and lifecycle retention settings. This lets the company define daily, weekly, and monthly retention centrally for supported EC2, EFS, and RDS resources. AWS Backup also provides centralized monitoring of backup jobs and job status through the AWS Backup console and APIs.

Configure an AWS backup plan to copy backups to another Region. is correct because a backup rule can include a copy action to a destination backup vault in another AWS Region. AWS Backup manages the cross-Region copy as part of the policy-driven backup workflow, avoiding custom replication automation.

Configure an Amazon EventBridge rule for AWS Backup Backup Job State Change events with a FAILED status, and set an Amazon SNS topic as the target. is correct because AWS Backup publishes job-state events to EventBridge. Filtering for failed backup jobs and routing matching events to SNS provides prompt notifications while remaining fully managed and applicable to protected resource types. See AWS Backup plans and rules and AWS Backup events in EventBridge.

QUESTION NO: 83

A company needs to migrate 500 TB of archived data from an on-premises data center to Amazon S3. The company's network connection cannot transfer the initial data set within the required 30-day migration window. After the initial transfer, the company needs to copy changed files over the network until cutover. The company requires data encryption and integrity validation.

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

A. Use an AWS Snowball Edge device to perform the initial bulk transfer of the archived data to Amazon S3.

B. Use AWS DataSync to perform recurring incremental transfers of changed files from the on-premises storage to Amazon S3 until cutover.

C. Use Amazon S3 Transfer Acceleration to transfer the initial 500 TB data set over the existing internet connection.

D. Use AWS Application Migration Service to replicate the archived files from the on-premises storage to Amazon S3.

E. Use Amazon FSx File Gateway to transfer the initial data set and synchronize changed files to Amazon S3.

ANSWER: A B

Explanation:

AWS Snowball Edge is suitable for transferring large initial data sets when available network bandwidth cannot meet the required migration timeline. Data copied to the device is encrypted, and AWS transfers the device data into the target Amazon S3 bucket after the device is returned.

AWS DataSync can perform repeated incremental transfers from the on-premises storage to Amazon S3 after the initial bulk copy. DataSync encrypts data in transit and supports verification options for validating transferred data. This combination minimizes the time required for the initial migration while maintaining synchronization until the final cutover. See AWS Snowball Edge and What is AWS DataSync?.

QUESTION NO: 84

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 Single Sign-On (AWS SSO) to use AWS Directory Service with an AD Connector that connects to the on-premises Active Directory. Grant access to the AWS accounts by assigning Active Directory groups to AWS SSO permission sets.

ANSWER: E

Explanation:

Configure AWS Single Sign-On (AWS SSO) to use AWS Directory Service with an AD Connector that connects to the on-premises Active Directory. Grant access to the AWS accounts by assigning Active Directory groups to AWS SSO permission sets. This solution preserves the company’s existing Active Directory as the authoritative, centralized location for users and groups. AD Connector is a directory gateway that lets IAM Identity Center, formerly AWS SSO, authenticate users against the on-premises Microsoft Active Directory without copying users or deploying separate IAM users for each AWS account. The existing Site-to-Site VPN provides the private network connectivity required between AWS and the on-premises domain controllers.

IAM Identity Center integrates with AWS Organizations to centrally administer access across member accounts. Administrators assign Active Directory users or groups to AWS accounts and permission sets. A permission set defines the permissions that IAM Identity Center provisions into each assigned account as an IAM role. Consequently, access can be granted and controlled according to the user’s Active Directory group membership and organizational job role, while identity lifecycle management remains in the existing directory. This meets the requirements for a single identity source and centralized, group-based access governance across the organization. For implementation details, see Manage your identity source in IAM Identity Center and IAM Identity Center permission sets.

QUESTION NO: 85

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

environment. The third account hosts a production environment. The application runs on Amazon ECS with the AWS Fargate launch type. The application stores data in an Amazon Aurora MySQL database.

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

The company wants to optimize costs for the application.

Which solution will meet this requirement?

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

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

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

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

ANSWER: B

Explanation:

In the organization management account, purchase Reserved Instances for a 1-year term for the production environment database. Purchase a 1-year Compute Savings Plan that will meet the compute usage demand. This approach aligns the commitments with the application's predictable, continuous production usage while avoiding an unnecessarily long commitment. Aurora MySQL supports Reserved DB Instances, which provide a billing discount in exchange for a one- or three-year term. Because the production database operates continuously and its demand is expected to remain stable for the next year, a one-year reservation is an appropriate commitment. AWS Organizations consolidated billing can share eligible reservation discounts across accounts, allowing the organization to centrally manage this discounted production database capacity. The ECS tasks use the Fargate launch type, and Compute Savings Plans apply automatically to eligible AWS Fargate usage, as well as Amazon EC2 and AWS Lambda. A one-year Compute Savings Plan can therefore cover the predictable baseline of the organization's Fargate spend while retaining flexibility for changing non-production usage. The development and testing environments have substantial scheduled downtime, so their intermittent database usage should not drive database reservation commitments. This combination targets the stable workload with discounted commitments and maintains flexibility where usage is not continuous. See Amazon RDS Reserved DB Instances and AWS Savings Plans documentation.

QUESTION NO: 86

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 C E

Explanation:

Creating an Amazon EventBridge (Amazon CloudWatch Events) rule that matches the `AWS API Call via CloudTrail` detail type and the `CreateUser` event name provides the event-driven trigger for this workflow. IAM `CreateUser` is a CloudTrail management event, and EventBridge can use CloudTrail-delivered API activity to initiate remediation immediately after a user is created. The EventBridge rule can run an Amazon ECS task on AWS Fargate. That task can use an IAM role with narrowly scoped permissions to remove the new user's effective access, such as deleting access keys and login profiles and removing permissions or group memberships according to the company's access-removal procedure. Fargate is appropriate because the remediation logic can be packaged as a container and run on demand without managing servers. The workflow can also publish to an Amazon SNS topic so that subscribed security-team endpoints receive the approval notification after the automated remediation action. SNS supports notification delivery to endpoints such as email, HTTPS, and AWS Lambda. For implementation details, see EventBridge events from CloudTrail and Amazon EventBridge targets.

QUESTION NO: 87

A company needs to migrate a 2 TB MySQL database from an on-premises data center to an Amazon Aurora cluster. The database receives hundreds of updates every minute. The on-premises database server is not accessible through the internet.

The migration solution must ensure that no data is lost between the start of migration and cutover. The migration must begin as soon as possible and must minimize downtime.

Which solution will meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the on-premises data center and the VPC that hosts the Aurora duster. Create a dump of the on-premises database by using mysqldump. Upload the dump to Amazon S3 by using multipart upload. Use an Amazon EC2 instance with appropriate permissions to import the dump to the Aurora cluster.

B. Create an AWS Site-to-Site VPN connection between the on-premises data center and the VPC that hosts the Aurora cluster. Specify the on-premises database as the source endpoint in AWS DMS. Specify the Aurora duster as the target endpoint. Configure a DMS task with ongoing replication.

C. Set up an AWS Direct Connect connection between the on-premises data center and the VPC that hosts the Aurora duster. Create a dump of the on-premises database by using mysqldump. Upload the dump to Amazon S3 by using multipart upload. Use an Amazon EC2 instance with appropriate permissions to import the dump to the Aurora cluster. Set up replication between the data center and the Aurora cluster.

D. Set up an AWS Direct Connect connection between the on-premises data center and the VPC that hosts the Aurora cluster. Specify the on-premises database as the source endpoint in AWS DMS. Specify the Aurora duster as the target endpoint Configure a DMS task with ongoing replication.

ANSWER: B

Explanation:

Create an AWS Site-to-Site VPN connection between the on-premises data center and the VPC that hosts the Aurora cluster. Specify the on-premises database as the source endpoint in AWS DMS. Specify the Aurora duster as the target endpoint. Configure a DMS task with ongoing replication. This approach provides private network connectivity promptly, allowing AWS Database Migration Service (AWS DMS) to access the on-premises MySQL source even though it is not exposed to the public internet. A Site-to-Site VPN can be established more quickly than a dedicated connectivity deployment, which supports the requirement to begin the migration as soon as possible.

AWS DMS can perform an initial full load of the existing 2 TB database and then continue with change data capture (CDC), replicating inserts, updates, and deletes from the MySQL binary logs to Aurora. Ongoing replication keeps the Aurora target synchronized while the initial load completes and while the application continues to write to the source. During cutover, the company can pause application writes, allow DMS to apply the final outstanding changes, validate that replication lag is cleared, and redirect the application to Aurora. This produces only a short write interruption and ensures changes made after the migration began are retained. See AWS DMS change data capture documentation and AWS Site-to-Site VPN documentation.

QUESTION NO: 88

A video processing company wants to build a machine learning (ML) model by using 600 TB of compressed data that is stored as thousands of files in the company's on-premises network attached storage system. The company does not have the necessary compute resources on premises for ML experiments and wants to use AWS. The company needs to complete the data transfer to AWS within 3 weeks. The data transfer will be a one-time transfer. The data must be encrypted in transit. The measured upload speed of the company's internet connection is 100 Mbps. and multiple departments share the connection. Which solution will meet these requirements MOST cost-effectively?

A. Order several AWS Snowball Edge Storage Optimized devices by using the AWS Management Console. Configure the devices with a destination S3 bucket. Copy the data to the devices. Ship the devices back to AWS.

B. Set up a 10 Gbps AWS Direct Connect connection between the company location and the nearest AWS Region. Transfer the data over a VPN connection into the Region to store the data in Amazon S3.

C. Create a VPN connection between the on-premises network attached storage and the nearest AWS Region. Transfer the data over the VPN connection.

D. Deploy an AWS Storage Gateway file gateway on premises. Configure the file gateway with a destination S3 bucket. Copy the data to the file gateway.

ANSWER: A

Explanation:

Order several AWS Snowball Edge Storage Optimized devices by using the AWS Management Console. Configure the devices with a destination S3 bucket. Copy the data to the devices. Ship the devices back to AWS. This is the most cost-effective approach because Snowball Edge is intended for large-scale, one-time offline migrations when available network capacity cannot meet the required transfer schedule. Even assuming the entire 100 Mbps internet upload bandwidth were continuously available, transferring 600 TB would take well over a year; shared departmental use would make the practical duration longer. Multiple Storage Optimized devices provide sufficient aggregate capacity to load the data locally from the NAS and physically transport it to AWS for import into Amazon S3 within the three-week requirement. Snowball Edge provides end-to-end security controls. Data stored on a device is encrypted with AWS Key Management Service (AWS KMS) keys, and communication used to manage and transfer data to the device is protected by encryption in transit. This avoids the recurring expense and provisioning lead time of high-bandwidth dedicated connectivity for a migration that will occur only once. AWS documents Snowball Edge as a solution for moving terabytes to petabytes of data when network transfer is impractical. See the AWS Snowball Edge Developer Guide and AWS Snow Family security documentation.

QUESTION NO: 89

A large company is migrating its entire IT portfolio to AWS. Each business unit in the company has a standalone AWS account that supports both development and test environments. New accounts to support production workloads will be needed soon. The finance department requires a centralized method for payment but must maintain visibility into each group's spending to allocate costs. The security team requires a centralized mechanism to control IAM usage in all the company's accounts. What combination of the following options meets the company's needs with the LEAST effort? (Choose two.)

A. Use a collection of parameterized AWS CloudFormation templates defining common IAM permissions that are launched into each account. Require all new and existing accounts to launch the appropriate stacks to enforce the least privilege model.

B. Use AWS Organizations to create a new organization from a chosen payer account and define an organizational unit hierarchy. Invite the existing accounts to join the organization and create new accounts using Organizations.

C. Require each business unit to use its own AWS accounts. Tag each AWS account appropriately and enable Cost Explorer to administer chargebacks.

D. Enable all features of AWS Organizations and establish appropriate service control policies that filter IAM permissions for sub-accounts.

E. Consolidate all of the company's AWS accounts into a single AWS account. Use tags for billing purposes and the IAM's Access Advisor feature to enforce the least privilege model.

ANSWER: B D

Explanation:

Use AWS Organizations to create a new organization from a chosen payer account and define an organizational unit hierarchy. Invite the existing accounts to join the organization and create new accounts using Organizations. This establishes a multi-account landing structure with a management account that centrally pays AWS charges through consolidated billing. Each member account remains separately identifiable in AWS billing and cost-management data, allowing finance to analyze, allocate, and charge back costs by business unit, environment, or account while avoiding separate payment administration. Organizations also streamlines the provisioning of the forthcoming production accounts and places accounts into organizational units that reflect the company's governance model.

Enable all features of AWS Organizations and establish appropriate service control policies that filter IAM permissions for sub-accounts. All features must be enabled to use service control policies (SCPs). SCPs provide centrally managed preventive guardrails by setting the maximum permissions that IAM principals may receive in accounts attached to the relevant organization root or organizational unit. The security team can therefore consistently restrict IAM-related and other AWS API actions across both existing member accounts and newly created production accounts. This combines centralized payment, granular cost visibility, scalable account governance, and centralized security controls with native AWS services and minimal operational effort. See AWS Organizations documentation and AWS service control policies documentation.

QUESTION NO: 90

A company is migrating an application from on-premises infrastructure to the AWS Cloud. During migration design meetings, the company expressed concerns about the availability and recovery options for its legacy Windows file server. The file server contains sensitive business-critical data that cannot be recreated in the event of data corruption or data loss. According to compliance requirements, the data must not travel across the public internet. The company wants to move to AWS managed services where possible. The company decides to store the data in an Amazon FSx for Windows File Server file system. A solutions architect must design a solution that copies the data to another AWS Region for disaster recovery (DR) purposes. Which solution will meet these requirements?

A. Create a destination Amazon S3 bucket in the DR Region. Establish connectivity between the FSx for Windows File Server file system in the primary Region and the S3 bucket in the DR Region by using Amazon FSx File Gateway. Configure the S3 bucket as a continuous backup source in FSx File Gateway.

B. Create an FSx for Windows File Server file system in the DR Region. Establish connectivity between the VPC the primary Region and the VPC in the DR Region by using AWS Site-to-Site VPN. Configure AWS DataSync to communicate by using VPN endpoints.

C. Create an FSx for Windows File Server file system in the DR Region. Establish connectivity between the VPC in the primary Region and the VPC in the DR Region by using VPC peering. Configure AWS DataSync to communicate by using interface VPC endpoints with AWS PrivateLink.

D. Create an FSx for Windows File Server file system in the DR Region. Establish connectivity between the VPC in the primary Region and the VPC in the DR Region by using AWS Transit Gateway in each Region. Use AWS Transfer Family to copy files between the FSx for Windows File Server file system in the primary Region and the FSx for Windows File Server file system in the DR Region over the private AWS backbone network.

ANSWER: C

Explanation:

Create an FSx for Windows File Server file system in the DR Region. Establish connectivity between the VPC in the primary Region and the VPC in the DR Region by using VPC peering. Configure AWS DataSync to communicate by using interface VPC endpoints with AWS PrivateLink. This design uses Amazon FSx for Windows File Server as managed file storage in both the primary and recovery Regions, while AWS DataSync provides managed, repeatable data-transfer tasks between the file systems. DataSync supports Amazon FSx for Windows File Server locations and can be used for ongoing synchronization that maintains a DR copy of the business-critical data.

Inter-Region VPC peering provides private IP connectivity between the VPCs across the AWS global network. The FSx file systems can therefore be reached through private routing rather than through public-facing paths. Additionally, an interface VPC endpoint for DataSync, powered by AWS PrivateLink, keeps communication between DataSync resources and the DataSync service private within AWS networking. This satisfies the requirement that data not traverse the public internet while avoiding self-managed file replication infrastructure. The recovery Region can retain a synchronized copy that is

available for DR activation if the primary Region or file system becomes unavailable. See the AWS DataSync documentation for FSx for Windows File Server locations and the AWS DataSync interface VPC endpoint documentation.

QUESTION NO: 91

A company is deploying a distributed in-memory database on a fleet of Amazon EC2 instances. The fleet consists of a primary node and eight worker nodes. The primary node is responsible for monitoring cluster health, accepting user requests, distributing user requests to worker nodes and sending an aggregate response back to a client. Worker nodes communicate with each other to replicate data partitions.

The company requires the lowest possible networking latency to achieve maximum performance.

Which solution will meet these requirements?

A. Launch memory optimized EC2 instances in a partition placement group

B. Launch compute optimized EC2 instances in a partition placement group

C. Launch memory optimized EC2 instances in a cluster placement group

D. Launch compute optimized EC2 instances in a spread placement group.

ANSWER: C

Explanation:

Launch memory optimized EC2 instances in a cluster placement group is the appropriate solution because this workload combines a large-memory requirement with extremely latency-sensitive, high-volume communication between nodes. Memory-optimized EC2 instance families are designed for workloads that process large datasets in memory, including in-memory databases and distributed caches. A cluster placement group places participating instances physically close together within a single Availability Zone. This placement is intended for tightly coupled workloads that need the lowest possible network latency, high network throughput, and high packet-per-second performance. Those characteristics directly support the frequent east-west traffic required for worker-node replication, as well as the primary node's request distribution and response aggregation. AWS recommends using supported instance types with enhanced networking when using cluster placement groups to obtain the available networking performance. The fleet is small enough to be placed in one cluster placement group, and launching the nodes together helps AWS allocate the necessary capacity. For additional resilience, the application should still replicate data appropriately and use recovery procedures, because a cluster placement group is scoped to one Availability Zone. See the AWS documentation on EC2 placement groups and the available memory-optimized Amazon EC2 instance types.

QUESTION NO: 92

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? (Choose 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:

Least-privilege connectivity is achieved by allowing only the application EC2 instances to initiate connections to the Aurora database listener, using security group references rather than broad IP CIDR ranges. The DB cluster security group should include an inbound rule whose source is the EC2 instances' security group and whose protocol and port match the Aurora engine's database listener port. This scopes database access to workloads that have the designated EC2 security group attached, including replacement or scaled instances without requiring IP-rule maintenance.

The EC2 instances' security group can also include an outbound rule with the DB cluster security group as its destination and the Aurora listener port as its port. This explicitly restricts egress from the application tier to the database tier and is appropriate when outbound access is being managed restrictively. Amazon VPC security groups are stateful: after a permitted connection is established, response traffic is automatically allowed. No separate rule for ephemeral response ports is needed. AWS documents security group references as a way to authorize traffic between resources based on their security groups, and Aurora access is controlled by VPC security group rules. See Amazon EC2 security group rules and Amazon Aurora security groups.

QUESTION NO: 93

A company is running several applications in the AWS Cloud. The applications are specific to separate business units in the company. The company is running the components of the applications in several AWS accounts that are in an organization in AWS Organizations. Every cloud resource in the company's organization has a tag that is named BusinessUnit. Every tag already has the appropriate value of the business unit name. The company needs to allocate its cloud costs to different business units. The company also needs to visualize the cloud costs for each business unit. Which solution will meet these requirements?

A. In the organization's management account, activate the existing BusinessUnit user-defined cost allocation tag. Also in the management account, create an Amazon S3 bucket and an AWS Cost and Usage Report (AWS CUR). Configure the S3 bucket as the destination for the AWS CUR. From the management account, query the AWS CUR data by using Amazon Athena. Use Amazon QuickSight for visualization.

B. In each member account, create a cost allocation tag that is named BusinessUnit. In the organization's management account, create an Amazon S3 bucket and an AWS Cost and Usage Report (AWS CUR). Configure the S3 bucket as the destination for the AWS CUR. Create an Amazon CloudWatch dashboard for visualization.

C. In the organization's management account, create a cost allocation tag that is named BusinessUnit. In each member account, create an Amazon S3 bucket and an AWS Cost and Usage Report (AWS CUR). Configure each S3 bucket as the destination for its respective AWS CUR. In the management account, create an Amazon CloudWatch dashboard for visualization.

D. In each member account, create a cost allocation tag that is named BusinessUnit. Also in each member account, create an Amazon S3 bucket and an AWS Cost and Usage Report (AWS CUR). Configure each S3 bucket as the destination for its respective AWS CUR. From the management account, query the AWS CUR data by using Amazon Athena. Use Amazon QuickSight for visualization.

ANSWER: A

Explanation:

In the organization’s management account, activate the existing user-defined `BusinessUnit` tag as a cost allocation tag, create a consolidated AWS Cost and Usage Report (AWS CUR) that delivers to Amazon S3, query the report with Amazon Athena, and use Amazon QuickSight for visualization. This uses the centralized billing and cost-management view available to the management account for all accounts in the AWS Organization. Activating the tag is the essential step: although the resources are already tagged, AWS does not include a user-defined tag in cost-management data until it has been activated as a cost allocation tag. Once activated, the tag becomes available as a cost dimension in the CUR, allowing usage and costs to be grouped by each business-unit value. A CUR provides the detailed, organization-wide line-item data needed for accurate allocation across accounts and services. Athena can query the CUR files in S3 with SQL, including aggregations by the `BusinessUnit` tag. QuickSight can then use the Athena query results to produce interactive dashboards and cost visualizations for stakeholders. See the AWS documentation for activating cost allocation tags and querying AWS CUR data with Athena.

QUESTION NO: 94

A company has an internet-facing application that runs on Amazon EC2 instances behind an Application Load Balancer. The application currently supports only IPv4. The company is running out of available private IPv4 addresses and needs to support IPv6 clients without changing the existing IPv4 connectivity.

Which solution will meet these requirements?

A. Associate an IPv6 CIDR block with the VPC. Assign IPv6 CIDR blocks to the application subnets. Change the Application Load Balancer IP address type to dualstack.

B. Create a NAT gateway in each Availability Zone. Configure the Application Load Balancer to use the NAT gateway IPv6 addresses.

C. Create an egress-only internet gateway. Configure the Application Load Balancer to use the egress-only internet gateway as its frontend endpoint.

D. Create a Network Load Balancer with an IPv6 target group. Register the existing IPv4 private addresses of the EC2 instances in the target group.

ANSWER: A

Explanation:

Associate an IPv6 CIDR block with the VPC, assign IPv6 CIDR blocks to the application subnets, and change the Application Load Balancer IP address type to dualstack. A dualstack Application Load Balancer receives both IPv4 and IPv6 addresses and can accept requests from IPv4 and IPv6 clients. The existing IPv4 configuration remains available while IPv6 capacity is added.

Each IPv6-enabled subnet receives a /64 IPv6 CIDR block from the VPC IPv6 allocation. Security groups and network ACLs must also be updated as necessary to permit IPv6 traffic. An internet gateway supports IPv6 internet connectivity without requiring a NAT gateway for outbound IPv6 traffic. See IPv6 addressing for Amazon VPC and Application Load Balancer IP address types.

QUESTION NO: 95

Question:

A company needs to copy backups of 40 RDS for MySQL databases from a production account to a central backup account within AWS Organizations. The databases usedefault AWS-managed KMS encryption keys. The backups must be stored in aWORM (Write Once Read Many)backup account.

What is the correct approach to enable cross-account backup?

A. Migrate the databases to customer-managed KMS keys, and use AWS Backup cross-account copy to a centrally owned Backup Vault Lock vault.

B. Share the default KMS keys with the central account and create backup vaults in the central account.

C. Use a Lambda function to decrypt and copy the snapshots to the central account.

D. Use a Lambda function to share and re-encrypt snapshots across accounts using the default KMS key.

ANSWER: A

Explanation:

Migrate the databases to customer-managed KMS keys, configure AWS Backup cross-account copy, and copy recovery points to a centrally owned Backup Vault Lock vault. This is required because encrypted Amazon RDS backups that use an AWS managed KMS key cannot be copied across AWS accounts: AWS managed keys cannot be shared or have their key policies changed to authorize use by another account. A customer-managed KMS key permits the required cross-account authorization through its key policy and grants, allowing AWS Backup to create encrypted copies in the destination account.

Because an RDS DB instance's encryption key cannot be changed in place, the practical migration path is to copy or restore the data using a snapshot encrypted under a customer-managed key, then protect the resulting DB instance with an AWS Backup plan. The central account can use AWS Backup cross-account copy within AWS Organizations to receive the backups. Applying AWS Backup Vault Lock to the destination vault enforces retention controls that provide WORM-style protection, preventing deletion or retention-period reduction during the lock's compliance period. See AWS Backup cross-account backup documentation and AWS Backup Vault Lock documentation.

QUESTION NO: 96

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 AccountB.

A solutions architect win deploy a two-net application In a new VPC To simplify the configuration, the db.example com CNAME record set tor 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 /eto/resolv.conf file

C. Create an authorization lo 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 m Account B Configure Route 53 replication between AWS accounts

E. Associate a new VPC in Account B with a hosted zone in Account

ANSWER: C E

Explanation:

A Route 53 private hosted zone only responds to DNS queries originating from VPCs that are explicitly associated with that hosted zone. Because the private hosted zone belongs to Account A and the new application VPC belongs to Account B, AWS requires a cross-account association workflow. First, the owner of the private hosted zone must create an authorization permitting the specified VPC in Account B to associate with the hosted zone. This authorization identifies the target VPC and its AWS Region. Then, using credentials in Account B, the VPC owner associates that VPC with the private hosted zone in Account A. Once the association is completed, workloads such as the application EC2 instance can use Route 53 Resolver through the VPC's Amazon-provided DNS service to resolve db.example.com to the CNAME record configured in the shared private hosted zone. This approach preserves centralized private DNS administration in Account A while allowing resources in Account B to resolve the necessary internal records. The relevant VPC must also have DNS support and DNS hostnames enabled as appropriate for normal Amazon-provided DNS resolution. AWS documents this exact cross-account process at Associating VPCs and private hosted zones across accounts and the authorization API at CreateVPCAssociationAuthorization.

QUESTION NO: 97

A company is using AWS Organizations with a multi-account architecture. The company's current security configuration for the account architecture includes SCPs, resource-based policies, identity-based policies, trust policies, and session policies. A solutions architect needs to allow an IAM user in Account A to assume a role in Account B. Which combination of steps must the solutions architect take to meet this requirement? (Choose three.)

A. Configure the SCP for Account A to allow the action.

B. Configure the resource-based policies to allow the action.

C. Configure the identity-based policy on the user in Account A to allow the action.

D. Configure the identity-based policy on the user in Account B to allow the action.

E. Configure the trust policy on the target role in Account B to allow the action.

F. Configure the session policy to allow the action and to be passed programmatically by the GetSessionToken API operation.

ANSWER: A C E

Explanation:

Cross-account role assumption requires authorization from both the calling principal's side and the target role's side, while remaining within AWS Organizations guardrails. "Configure the SCP for Account A to allow the action." is required in the sense that the SCPs effective for Account A must permit `sts:AssumeRole`; SCPs establish the maximum permissions that IAM identities in a member account can receive. An applicable explicit deny, or an allow-list SCP structure that does not allow the action, prevents the request regardless of IAM policy permissions.

"Configure the identity-based policy on the user in Account A to allow the action." grants the IAM user permission to call `sts:AssumeRole` against the specific role ARN in Account B. "Configure the trust policy on the target role in Account B to allow the action." establishes that the target role trusts the source IAM user, its account, or an appropriate principal in Account A. Together, these permissions form the required cross-account authorization path: the caller is authorized to make the STS request, and the target role authorizes that caller to assume it. AWS documents these caller-permission and role-trust requirements for cross-account role access, along with the permission-boundary effect of SCPs. See AWS IAM role switching permissions and AWS Organizations service control policies.

QUESTION NO: 98

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

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

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

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

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

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

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

ANSWER: B

Explanation:

Create an IAM OpenID Connect (OIDC) identity provider (IdP) in IAM. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRoleWithWebIdentity API call from the GitHub OIDC IdP. Update GitHub to assume the role for the pipeline. This is the lowest-overhead solution because GitHub Actions natively issues OIDC tokens for workflow jobs. The workflow requests an ID token from GitHub and presents it to AWS Security Token Service by using

`AssumeRoleWithWebIdentity`. AWS then returns automatically rotated, short-lived role credentials, eliminating the IAM user's long-lived access key and secret key from the pipeline.

The role can use the existing deployment permissions, and its trust policy can limit assumption using GitHub OIDC token claims, including the specific GitHub organization, repository, branch, tag, or deployment environment. This provides a narrowly scoped, workload identity-based authentication model without maintaining static credentials, custom federation infrastructure, or client certificates. AWS documents GitHub Actions OIDC as a direct federated identity integration with IAM, while GitHub provides an official AWS workflow configuration using the `id-token: write` permission. See AWS IAM OIDC identity providers and GitHub Actions: Configuring OpenID Connect in AWS.

QUESTION NO: 99

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

The company needs to automate this process.

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

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

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

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

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

ANSWER: B

Explanation:

Create a Service Quotas request template in the organization's management account, configure it with the required Amazon EC2 quota increase, and associate it with the organization. Service Quotas supports AWS Organizations templates specifically to standardize quota-increase requests across member accounts. After the template is associated with the organization, AWS Service Quotas automatically creates the specified quota-increase requests for existing member accounts and for accounts that join the organization later, including newly created accounts. This removes the need for anyone to sign in as a root user or to manually submit each request.

This is the lowest-operational-overhead approach because it is an AWS-managed, declarative organization-level feature. The company defines the approved EC2 quota and target region once, then AWS applies the request consistently as accounts are onboarded. The management account must have the required Service Quotas and AWS Organizations permissions, and quota requests remain subject to AWS review where the quota is adjustable. The template therefore automates request submission, while preserving AWS's normal quota approval process. For implementation details, see the AWS Service Quotas documentation for AWS Organizations templates and the AWS Organizations account creation documentation.

QUESTION NO: 100

A company is planning to migrate an Amazon RDS for Oracle database to an RDS for PostgreSQL DB instance in another AWS account. A solutions architect needs to design a migration strategy that will require no downtime and that will minimize the amount of time necessary to complete the migration. The migration strategy must replicate all existing data and any new data that is created during the migration The target database must be identical to the source database at completion of the migration process

All applications currently use an Amazon Route 53 CNAME record as their endpoint for communication with the RDS for Oracle DB instance The RDS for Oracle DB instance is in a private subnet.

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

A. Create a new RDS for PostgreSQL DB instance in the target account Use the AWS Schema Conversion Tool (AWS SCT) to migrate the database schema from the source database to the target database

B. Use the AWS Schema Conversion Tool (AWS SCT) to create a new RDS for PostgreSQL DB instance in the target account with the schema and initial data from thesource database

C. Configure VPC peering between the VPCs in the two AWS accounts to provide connectivity to both DB instances from the target account. Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account.

D. Temporarily allow the source DB instance to be publicly accessible to provide connectivity from the VPC in the target account Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account.

E. Use AWS Database Migration Service (AWS DMS) in the target account to perform a full load plus change data capture (CDC) migration from the source database to the target database When the migration is complete, change the CNAME record to point to the target DB instance endpoint

F. Use AWS Database Migration Service (AWS DMS) in the target account to perform a change data capture (CDC) migration from the source database to the target database When the migration is complete change the CNAME record to pointto the target DB instance endpoint.

ANSWER: A C E

Explanation:

Creating an RDS for PostgreSQL DB instance and using AWS Schema Conversion Tool (AWS SCT) to migrate the database schema establishes a compatible PostgreSQL target for this heterogeneous Oracle-to-PostgreSQL migration. AWS SCT assesses and converts supported Oracle schema objects and produces the converted schema for the target database. Private connectivity is needed because the Oracle source is not publicly accessible. VPC peering between the two accounts' VPCs, along with routes and database-port security-group permissions, enables the migration components to communicate with both databases over private IP addresses. AWS DMS can then run a full load plus change data capture migration. The full-load phase copies the existing source data, while CDC continuously applies inserts, updates, and deletes that occur after the full load begins. This minimizes the final synchronization window and allows the PostgreSQL target to catch up with the Oracle source before the endpoint change. Once replication latency is zero or acceptably low and final validation is complete, changing the Route 53 CNAME to the new RDS endpoint directs applications to PostgreSQL. The CNAME abstraction avoids application connection-string changes and supports a controlled cutover. See the AWS DMS documentation for full load and CDC tasks and the AWS Schema Conversion Tool User Guide.

QUESTION NO: 101

To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network. The security team mandates that no traffic accessing this data should traverse the public internet.

How should the solutions architect design a highly available solution that meets the requirements and is cost-effective?

A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use. Use the company WAN lo send traffic over to the headquarters and then to the respective DX connection to access the data.

B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection. Use inter-region VPC peering to access the data in other AWS Regions.

C. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection. Use an AWS transit VPC solution to access data in other AWS Regions.

D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection. Use Direct Connect Gateway to access data in other AWS Regions.

ANSWER: D

Explanation:

Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection. Use Direct Connect Gateway to access data in other AWS Regions. This design provides dedicated private connectivity between the company WAN and AWS, so traffic does not traverse the public internet. For availability, the two Direct Connect connections provide redundant paths; in a production implementation, AWS recommends using connections at separate Direct Connect locations to reduce exposure to a location or device failure.

A Direct Connect gateway is the appropriate scalable mechanism for connecting a Direct Connect virtual interface to resources in multiple AWS Regions. The company can associate the Direct Connect gateway with virtual private gateways or transit gateways that provide access to the required VPCs, including VPCs outside the Region where the Direct Connect connections terminate. This avoids deploying and paying for separate Direct Connect connectivity in every Region, while maintaining private AWS network connectivity for inter-Region access. Route configuration and allowed prefixes should be controlled carefully through the Direct Connect gateway associations. See the AWS Direct Connect gateway documentation and AWS Direct Connect User Guide.

QUESTION NO: 102

A company uses an Amazon Aurora PostgreSQL DB cluster for applications in a single AWS Region. The company's database team must monitor all data activity on all the databases. Which solution will achieve this goal?

A. Set up an AWS Database Migration Service (AWS DMS) change data capture (CDC) task. Specify the Aurora DB cluster as the source. Specify Amazon Kinesis Data Firehose as the target. Use Kinesis Data Firehose to upload the data into an Amazon OpenSearch Service cluster for further analysis.

B. Start a database activity stream on the Aurora DB cluster to capture the activity stream in Amazon EventBridge. Define an AWS Lambda function as a target for EventBridge. Program the Lambda function to decrypt the messages from EventBridge and to publish all database activity to Amazon S3 for further analysis.

C. Start a database activity stream on the Aurora DB cluster to push the activity stream to an Amazon Kinesis data stream. Configure Amazon Kinesis Data Firehose to consume the Kinesis data stream and to deliver the data to Amazon S3 for further analysis.

D. Set up an AWS Database Migration Service (AWS DMS) change data capture (CDC) task. Specify the Aurora DB cluster as the source. Specify Amazon Kinesis Data Firehose as the target. Use Kinesis Data Firehose to upload the data into an Amazon Redshift cluster. Run queries on the Amazon Redshift data to determine database activities on the Aurora database.

ANSWER: C

Explanation:

Start a database activity stream on the Aurora DB cluster to push the activity stream to an Amazon Kinesis data stream. Configure Amazon Kinesis Data Firehose to consume the Kinesis data stream and to deliver the data to Amazon S3 for further analysis. This uses Aurora Database Activity Streams, the Aurora-native auditing feature for capturing database activity at the DB-cluster level. The stream provides near-real-time audit records for activity across the databases in the Aurora PostgreSQL cluster, enabling the database team to centrally monitor and retain that activity without modifying individual applications or implementing database triggers.

When the activity stream is started, Aurora publishes the audit records to an Amazon Kinesis Data Streams stream. The records are protected with the AWS KMS key configured for the activity stream, supporting controlled access to potentially sensitive audit content. Amazon Data Firehose can consume records from that Kinesis data stream and reliably deliver them to Amazon S3. S3 then provides durable, low-cost retention and serves as a destination for downstream auditing, querying, and security-analysis workflows. Consumers that need to inspect the record contents must have appropriate KMS permissions to decrypt them. See Aurora Database Activity Streams and Using Kinesis Data Streams as an Amazon Data Firehose source.

QUESTION NO: 103

A company has developed an application that is running Windows Server on VMware vSphere VMs that the company hosts on premises The application data is stored in a proprietary format that must be read through the application The company manually provisioned the servers and the application

As part of its disaster recovery plan, the company wants the ability to host its application on AWS temporarily if the company ' s on-premises environment becomes unavailable The company wants the application to return to on-premises hosting after a disaster recovery event is complete The RPO is 5 minutes.

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

A. Configure AWS DataSync Replicate the data to Amazon Elastic Block Store (Amazon EBS) volumes When the on-premises environment is unavailable, use AWS Cloud Format ion templates to provision Amazon EC2 instances and attach the EBS volumes

B. Configure AWS Elastic Disaster Recovery Replicate the data to replication Amazon EC2 instances that are attached to Amazon Elastic Block Store (Amazon EBS) volumes When the on-premises environment is unavailable use Elastic Disaster Recovery to launch EC2 instances that use the replicated volumes

C. Provision an AWS Storage Gateway file gateway. Replicate the data to an Amazon S3 bucket When the on-premises environment is unavailable, use AWS Backup to restore the data to Amazon Elastic Block Store (Amazon EBS) volumes and launch Amazon EC2 instances from these EBS volumes

D. Provision an Amazon FSx for Windows File Server file system on AWS Replicate the data to the file system When the on-premises environment is unavailable, use AWS Cloud Format ion templates to provision Amazon EC2 instances and use AWS CloudFormation Init commands to mount the Amazon FSx file shares

ANSWER: B

Explanation:

Configure AWS Elastic Disaster Recovery Replicate the data to replication Amazon EC2 instances that are attached to Amazon Elastic Block Store (Amazon EBS) volumes When the on-premises environment is unavailable use Elastic Disaster Recovery to launch EC2 instances that use the replicated volumes is the best fit because AWS Elastic Disaster Recovery is purpose-built to recover physical, virtual, and cloud-based servers into AWS with minimal operational management. An agent installed on each source Windows Server continuously performs block-level replication into a low-cost AWS staging area. This preserves the server’s disks, installed application, configuration, and proprietary application data together, so recovery instances can run the same application stack required to read the data.

Elastic Disaster Recovery supports continuous replication with an RPO measured in seconds under normal operating conditions, which satisfies the required 5-minute RPO. During an outage, the company can launch recovery EC2 instances from selected recovery points using predefined launch settings. After the on-premises environment is restored, Elastic Disaster Recovery supports failback by replicating changes from AWS recovery instances back to the original source infrastructure. This directly supports temporary AWS hosting followed by a return to on-premises operation, without requiring the company to redesign the application, convert its proprietary data format, or build and operate custom replication and restoration workflows. For details, see the AWS Elastic Disaster Recovery User Guide and the AWS disaster recovery overview.

QUESTION NO: 104

A company runs a highly available data collection application on Amazon EC2 in the eu-north-1 Region. The application collects data from end-user devices and writes records to an Amazon Kinesis data stream and a set of AWS Lambda functions that process the records. The company persists the output of the record processing to an Amazon S3 bucket in eu-north-1. The company uses the data in the S3 bucket as a data source for Amazon Athena.

The company wants to increase its global presence. A solutions architect must launch the data collection capabilities in the sa-east-1 and ap-northeast-1 Regions. The solutions architect deploys the application, the Kinesis data stream, and the Lambda functions in the two new Regions. The solutions architect keeps the S3 bucket in eu-north-1 to meet a requirement to centralize the data analysis.

During testing of the new setup, the solutions architect notices a significant lag on the arrival of data from the new Regions to the S3 bucket.

Which solution will improve this lag time the MOST?

A. In each of the two new Regions, set up the Lambda functions to run in a VPC.

B. Turn on S3 Transfer Acceleration on the S3 bucket in eu-north-1. Change the application to use the new S3 accelerated endpoint when the application uploads data to the S3 bucket.

C. Create an S3 bucket in each of the two new Regions. Set the application in each new Region to upload to its respective S3 bucket. Set up S3 Cross-Region Replication to replicate data to the S3 bucket in eu-north-1.

D. Increase the memory requirements of the Lambda functions to ensure that they have multiple cores available. Use the multipart upload feature when the application uploads data to Amazon S3 from Lambda.

ANSWER: B

Explanation:

Turn on S3 Transfer Acceleration on the S3 bucket in eu-north-1. Change the application to use the new S3 accelerated endpoint when the application uploads data to the S3 bucket. This directly addresses the latency-sensitive, cross-Region upload path from South America and Asia Pacific to the centrally located bucket in Europe. S3 Transfer Acceleration uses Amazon CloudFront edge locations to receive uploads close to the client or workload, then transfers the data over AWS's optimized global network to the target S3 bucket. This can substantially improve transfer speed for uploads that traverse long geographic distances, while preserving the required single, centralized S3 data source for Athena.

The workload must use the bucket's acceleration endpoint for the benefit to apply. Transfer Acceleration does not require changing the bucket's Region, creating intermediate storage locations, or adding an asynchronous replication stage; processed records can continue to be written directly to the central bucket. AWS documents that Transfer Acceleration is intended to enable fast, easy, and secure transfers over long distances between clients and an S3 bucket. See the Amazon S3 Transfer Acceleration documentation and the Transfer Acceleration endpoint examples.

QUESTION NO: 105

A company wants to migrate its website to AWS. The website uses containers that are deployed in an on-premises, self-managed Kubernetes cluster. All data for the website is stored in an on-premises PostgreSQL database.

The company has decided to migrate the on-premises Kubernetes cluster to an Amazon EKS cluster. The EKS cluster will use EKS managed node groups with a static number of nodes. The company will also migrate the on-premises database to an Amazon RDS for PostgreSQL database.

A solutions architect needs to estimate the total cost of ownership (TCO) for this workload before the migration.

Which solution will provide the required TCO information?

A. Request access to Migration Evaluator. Run the Migration Evaluator Collector and import the data. Configure a scenario. Export a Quick Insights report from Migration Evaluator.

B. Launch AWS DMS for the on-premises database. Generate an assessment report. Create an estimate in AWS Pricing Calculator for the costs of the EKS migration.

C. Initialize AWS Application Migration Service. Add the on-premises servers as source servers. Launch a test instance. Output a TCO report from Application Migration Service.

D. Access the AWS Cloud Economics Center webpage to assess the AWS Cloud Value Framework. Create an AWS Cost and Usage report from the Cloud Value Framework.

ANSWER: A

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. AWS Migration Evaluator is specifically intended to develop a data-driven migration business case, including projected AWS costs and total cost of ownership before a migration begins. The Collector gathers information about the existing on-premises environment, including inventory and resource-utilization data. This gives the assessment an evidence-based baseline rather than relying only on assumed resource sizes.

After the collected data is imported, the solutions architect can configure a target scenario that models the planned AWS environment, such as the required compute capacity for the static Amazon EKS managed node group and the target Amazon RDS for PostgreSQL deployment. Migration Evaluator applies AWS pricing and can model costs over time, enabling a comparison between the current environment and the proposed AWS infrastructure. Its Quick Insights report presents the resulting cost and savings analysis in a concise form suitable for stakeholders assessing the migration decision. This directly meets the requirement for a pre-migration TCO estimate. See the AWS Migration Evaluator User Guide and AWS Migration Evaluator product page.

QUESTION NO: 106

A company is planning a one-time migration of an on-premises MySQL database to Amazon Aurora MySQL in the us-east-1 Region. The company ' s current internet connection has limited bandwidth. The on-premises MySQL database is 60 TB in size The company estimates that it will take a month to transfer the data to AWS over the current internet connection.

The company needs a migration solution that will migrate the database more quickly

Which solution will migrate the database in the LEAST amount of time?

A. Request a 1 Gbps AWS Direct Connect connection between the on-premises data center and AWS Use AWS Database Migration Service (AWS DMS) to migrate the on-premises MySQL database to Aurora MySQL.

B. Use AWS DataSync with the current internet connection to accelerate the data transfer between the on-premises data center and AWS Use AWS Application Migration Service to migrate the on-premises MySQL database to Aurora MySQL.

C. Order an AWS Snowball Edge device. Load an export of the database data onto the device by using its Amazon S3-compatible interface, then use AWS Database Migration Service (AWS DMS) to migrate the data from Amazon S3 to Aurora MySQL.

D. Order an AWS Snowball Device Load the data into an Amazon S3 bucket by usingthe S3 Adapter for Snowball Use AWS Application Migration Service to migrate the data from Amazon S3 to Aurora MySQL.

ANSWER: C

Explanation:

Order an AWS Snowball Edge device, load an export of the database data onto the device through its Amazon S3-compatible interface, and use AWS DMS to load the data from Amazon S3 into Aurora MySQL. This approach is designed for a large, one-time transfer when available network bandwidth makes online migration impractical. A Snowball Edge Storage Optimized device provides sufficient usable capacity for a 60 TB dataset and allows the company to copy the exported database files locally at LAN speeds rather than uploading the full dataset through the constrained internet connection. After the device is returned, AWS transfers the data into the specified Amazon S3 bucket in the target Region.

For this migration pattern, AWS DMS can use Amazon S3 as a source endpoint and Aurora MySQL as the target. The database export must be prepared in an S3-source format that DMS supports, such as CSV files with the required table and column mapping. Because the workload is a one-time migration, an offline bulk transfer removes the month-long network-transfer bottleneck; DMS then performs the database load within AWS. AWS documents Snowball Edge as a service for moving large amounts of data into AWS when network transfer is constrained, and documents Amazon S3 source endpoints for DMS at AWS Snowball Edge documentation and AWS DMS S3 source documentation.

QUESTION NO: 107

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:

Set the viewer protocol policy to Redirect HTTP to HTTPS for CloudFront so a viewer who initially requests an HTTP URL receives an HTTP redirect response and is then connected over HTTPS. This provides a seamless experience for HTTP users while ensuring that normal website communication is encrypted between the viewer and the CloudFront distribution. Set the origin protocol policy to HTTPS Only for CloudFront so that CloudFront uses TLS when it forwards requests to the Application Load Balancer. This matches the existing ACM certificate configured on the ALB HTTPS listener and encrypts the CloudFront-to-origin connection. Create an S3 bucket policy that denies access to the S3 bucket if the aws:SecureTransport request is false. The `aws:SecureTransport` global condition key lets the bucket enforce TLS for requests that retrieve the website content, including requests made by the Lambda application. These controls provide HTTPS enforcement across the viewer, CloudFront origin, and S3 content-access paths, while the redirect policy specifically prevents HTTP users from receiving an access error. See CloudFront HTTPS viewer and origin protocol policies and Amazon S3 policy condition keys.

QUESTION NO: 108

A company hosts a VPN in an on-premises data center. Employees currently connect to the VPN to access files in their Windows home directories. Recently, there has been a large growth in the number of employees who work remotely. As a result, bandwidth usage for connections into the data center has begun to reach 100% during business hours. The company must design a solution on AWS that will support the growth of the company's remote workforce, reduce the bandwidth usage for connections into the data center, and reduce operational overhead. Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose two.)

A. Create an AWS Storage Gateway Volume Gateway. Mount a volume from the Volume Gateway to the on-premises file server.

B. Migrate the home directories to Amazon FSx for Windows File Server.

C. Migrate the home directories to Amazon FSx for Lustre.

D. Migrate remote users to AWS Client VPN.

E. Create an AWS Direct Connect connection from the on-premises data center to AWS.

ANSWER: B D

Explanation:

Migrate the home directories to Amazon FSx for Windows File Server. Amazon FSx for Windows File Server is a fully managed service designed to provide Windows-native file storage. It supports the SMB protocol, integration with Microsoft Active Directory, and Windows file system features such as NTFS permissions and access control lists. These capabilities make it suitable for hosting existing Windows home directories without requiring the company to operate, patch, scale, or maintain conventional Windows file servers. FSx also provides managed storage and availability features, reducing the operational burden as the workforce and its file-storage requirements grow. AWS details these capabilities in the Amazon FSx for Windows File Server documentation.

Migrate remote users to AWS Client VPN. AWS Client VPN is a managed remote-access VPN service that lets employees establish secure client connections directly to AWS resources. With home directories hosted on FSx in AWS, users access their files through AWS rather than first connecting to the on-premises VPN and consuming constrained inbound data-center bandwidth. The managed endpoint model can scale with remote-user demand while eliminating the need to build and

operate a self-managed remote-access VPN fleet. AWS Client VPN supports integration with directory-based authentication and authorization, which is useful for an existing Windows-oriented environment. See the AWS Client VPN documentation.

QUESTION NO: 109

A company wants to migrate its on-premises data center to the AWS Cloud. This includes thousands of virtualized Linux and Microsoft Windows servers, SAN storage, Java and PHP applications with MYSQL, and Oracle databases. There are many dependent services hosted either in the same data center or externally.

The technical documentation is incomplete and outdated. A solutions architect needs to understand the current environment and estimate the cloud resource costs after the migration.

Which tools or services should solutions architect use to plan the cloud migration? (Choose three.)

A. AWS Application Discovery Service

B. AWS SMS

C. AWS x-Ray

D. AWS Cloud Adoption Readiness Tool (CART)

E. Amazon Inspector

F. AWS Migration Hub

ANSWER: A D F

Explanation:

AWS Application Discovery Service is appropriate because it gathers information from the on-premises estate when existing documentation cannot be relied upon. Its agents and agentless collector can capture server inventory, configuration, utilization, and network-connection data. This discovery data enables the architect to identify application dependencies, organize servers into migration groups, and obtain the sizing inputs needed to estimate AWS resource requirements and costs. AWS Migration Hub provides the central migration-planning workspace for discovered servers and applications. It can ingest discovery data, help group resources into applications, and provide visibility into migration planning and progress across supported AWS migration tools. AWS Cloud Adoption Readiness Tool (CART) is also applicable to the planning phase represented by this question because it assesses organizational cloud-adoption readiness across the business, people, process, platform, operations, and security perspectives. The resulting readiness assessment identifies migration-program gaps and recommended actions alongside the technical discovery work. Together, these services support a structured assessment of both the existing technical environment and the organizational preparation required before migration execution. AWS describes discovery capabilities in the Application Discovery Service User Guide and migration application tracking in the AWS Migration Hub User Guide.

QUESTION NO: 110

A retail company has structured its AWS accounts to be part of an organization in AWS Organizations. The company has set up consolidated billing and has mapped its departments to the following OUs: Finance, Sales, Human Resources (HR), Marketing, and Operations. Each OU has multiple AWS accounts, one for each environment within a department. These environments are development, test, pre-production, and production. The HR department is releasing a new system that will launch in 3 months. In preparation, the HR department has purchased several Reserved Instances (RIs) in its production AWS account. The HR department will install the new application on this account. The HR department wants to make sure that other departments cannot share the RI discounts. Which solution will meet these requirements?

A. In the AWS Billing and Cost Management console for the HR department's production account turn off RI sharing.

B. Remove the HR department's production AWS account from the organization. Add the account 10 the consolidating billing configuration only.

C. In the AWS Billing and Cost Management console. use the organization's management account 10 turn off RI Sharing for the HR departments production AWS account.

D. Create an SCP in the organization to restrict access to the RIs. Apply the SCP to the OUs of the other departments.

ANSWER: C

Explanation:

In the AWS Billing and Cost Management console. use the organization's management account 10 turn off RI Sharing for the HR departments production AWS account. This is the appropriate solution because Reserved Instance discount sharing is managed centrally by the AWS Organizations management account for accounts that use consolidated billing. By default, eligible unused RI discounts can be shared across member accounts in an organization, allowing matching usage in other accounts to receive the discount. The management account can change this behavior for an individual member account by disabling its RI discount sharing preference.

After sharing is disabled for the HR production account, that account retains the ability to receive the RI discount for its own matching usage. Its unused RI discount capacity, however, is not shared with other organization accounts. This preserves the company's AWS Organizations structure and consolidated billing arrangement while ensuring that the RIs bought by HR are reserved for HR production-account usage. AWS documents that RI discount sharing preferences are configured through Billing and Cost Management by the management account and can be set on a per-account basis. See AWS consolidated billing and discount sharing documentation and Turning off Reserved Instance discount sharing.

QUESTION NO: 111

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 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? (Select TWO.)

A. Create an Amazon CloudFront distribution with the ALB as the origin. Add a custom header and random value to the CloudFront origin configuration. 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 its origin places the application behind AWS's globally distributed edge network. Configuring a secret custom origin header in CloudFront and an ALB listener rule that forwards requests only when that header and value are present helps ensure that requests reaching the application follow the intended CloudFront path. This reduces exposure to direct requests against the ALB and preserves the application's normal request-routing behavior during an attack.

Associating an AWS WAF web ACL with the CloudFront distribution provides inspection and filtering at edge locations before malicious requests consume ALB, ECS, or DynamoDB capacity. The web ACL can use appropriate AWS Managed Rules rule groups, along with rate-based rules when needed, to block common web exploits, malicious request patterns, and excessive request rates. CloudFront and AWS WAF are managed, elastic services, so they can absorb and filter traffic without requiring the company to pre-provision additional application containers or database caching capacity solely for an attack. Together, these controls protect the public entry point, reduce load on the application origin, and support continuity with minimal interruption. See CloudFront custom origin headers and associating an AWS WAF web ACL with CloudFront.

QUESTION NO: 112

An AWS partner company is building a service in AWS Organizations using Its organization named org. This service requires the partner company to have access to AWS resources in a customer account, which is in a separate organization named org2 The company must establish least privilege security access using an API or command line tool to the customer account

What is the MOST secure way to allow org1 to access resources h org2?

A. The customer should provide the partner company with their AWS account access keys to log in and perform the required tasks

B. The customer should create an IAM user and assign the required permissions to the IAM user The customer should then provide the credentials to the partner company to log In and perform the required tasks.

C. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM rote's Amazon Resource Name (ARN) when requesting access to perform the required tasks

D. The customer should create an IAM rote and assign the required permissions to the IAM rote. The partner company should then use the IAM rote's Amazon Resource Name (ARN). Including the external ID in the IAM role's trust pokey, when requesting access to perform the required tasks

ANSWER: D

Explanation:

The customer should create an IAM role with only the required permissions, and the partner should assume that role by using its ARN and the required external ID. This is the AWS-recommended cross-account pattern for granting a third party API or CLI access without distributing long-term credentials. The role's permissions policy is owned by the customer and can be scoped precisely to the required resources and actions, implementing least privilege. The role trust policy identifies the partner's AWS principal and requires the externally supplied value as a condition of assuming the role. When the partner calls AWS STS AssumeRole, AWS returns temporary credentials that are limited by the role policy and expire automatically. The external ID is particularly important for a partner that serves multiple customers. It mitigates the confused deputy risk by requiring a value unique to the customer relationship, so the partner cannot accidentally use its access path for one customer to obtain credentials for another customer. The customer should generate or agree on the external ID and configure it in the role trust policy; the partner supplies the matching value in its AssumeRole request. See AWS guidance on preventing the confused deputy problem and the AWS STS AssumeRole API.

QUESTION NO: 113

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 fetched from the closest Region. If game assess become unavailable in the closest Region, they should the fetched from the other Region.

What should a solutions architect do to meet these requirement?

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 tor each ALB. Create a Route 53 failover routing record pointing to the two ALBs. Set the Evaluate Target Health value 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 tor 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. This design directly addresses both routing requirements at the DNS layer. Route 53 latency-based routing evaluates latency measurements between the requester's resolver and the AWS Regions hosting the record targets, then returns the record for the Region expected to provide the lowest latency. As a result, players are directed to the Regional ALB that is closest in network-latency terms, rather than to a fixed primary Region.

The health checks and target-health evaluation provide regional availability protection. Route 53 can consider an ALB target unhealthy when its underlying load balancer target health indicates that it cannot serve traffic. If the lowest-latency Regional endpoint is unhealthy, Route 53 stops returning that endpoint and answers DNS queries with the healthy latency record for the other Region. This provides automatic failover while retaining normal low-latency routing. DNS clients and resolvers can continue using a previously cached response until its TTL expires, so an appropriately short TTL should be selected for the application's recovery objectives. See the AWS documentation for latency-based routing and Route 53 health checks and DNS failover.

QUESTION NO: 114

A new application is running on Amazon Elastic Container Service (Amazon ECS) with AWS Fargate The application uses an Amazon Aurora MySQL database The application and the database run m the same subnets of a VPC with distinct security groups that are configured.

The password (or the database is stored m AWS Secrets Manager and is passed to the application through the D8_PASSWORD environment variable The hostname of the database is passed to the application through the DB_HOST environment variable The application Is failing to access the database.

Which combination of actions should a solutions architect take to resolve this error? (Select THREE)

A. Ensure that the container has the environment variable with name "DB_PASSWORD" specified with a "ValueFrom" and the ARN of the secret

B. Ensure that the container has the environment variable with name *D8_PASSWORD" specified with a "ValueFrom" and the secret name of the secret.

C. Ensure that the Fargate service security group allows inbound network traffic from the Aurora MySQL database on the MySQL TCP port 3306.

D. Ensure that the Aurora MySQL database security group allows inbound network traffic from the Fargate service on the MySQL TCP port 3306.

E. Ensure that the container has the environment variable with name "D8_HOST" specified with the hostname of a DB instance endpoint.

F. Ensure that the container has the environment variable with name "DB_HOST" specified with the hostname of the OB duster endpoint.

ANSWER: A D F

Explanation:

The correct resolution is to inject the database credential into the container through the ECS task definition, permit the application-to-database network flow, and use the Aurora cluster connection endpoint. The container definition should define `DB_PASSWORD` as a secret environment variable whose `ValueFrom` value identifies the AWS Secrets Manager secret ARN. ECS retrieves the secret at task startup and exposes it to the running container without placing the password directly in the task definition. The task execution role must also have permission to retrieve the referenced secret.

Because the ECS task initiates a MySQL connection to Aurora, the Aurora DB security group must have an inbound rule for TCP port 3306 with the Fargate service or task security group as its source. Referencing the source security group permits this private VPC traffic regardless of the dynamically assigned task IP address. Finally, `DB_HOST` should contain the Aurora DB cluster endpoint hostname. The cluster endpoint is the primary endpoint for an Aurora cluster and directs connections appropriately to the cluster writer. See the Amazon ECS documentation for Secrets Manager environment variables and the Amazon Aurora endpoint documentation.

QUESTION NO: 115

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

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

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

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

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

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

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

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

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

ANSWER: A B

Explanation:

Create an identity policy that grants the user read and write access. Add a condition that specifies that the S3 paths must be prefixed with ${aws:username}. Apply the policy on the scientists ' IAM user group. This approach uses the IAM policy variable `${aws:username}` to dynamically scope S3 permissions to each requesting IAM user's own object prefix. A single group policy can therefore enforce individual folder isolation for all scientists, without requiring separately managed policies for each person. The policy should grant object permissions only for the username-specific prefix and can use the `s3:prefix` condition for appropriately scoped bucket listing. AWS documents IAM policy variables and their use in resource ARNs and condition values at IAM policy variables.

Configure a trail with AWS CloudTrail to capture all object-level events in the S3 bucket. Store the trail output in another S3 bucket. Use Amazon Athena to query the logs and generate reports. CloudTrail S3 data events record object-level API activity, including reads and writes, with the identity that made the request, the affected object, time, and source details. Delivering those logs to S3 and querying them with Athena creates a managed, serverless reporting process with low operational effort. S3 data-event logging is described in Logging Amazon S3 API calls using CloudTrail.

QUESTION NO: 116

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. This architecture uses managed, scalable services for each stage of the workload while retaining MQTT as the device communication protocol. AWS IoT Core provides a managed MQTT message broker that can securely support large fleets of connected devices without requiring the company to operate or recover broker servers. IoT rules can evaluate inbound messages and send matching records to an Amazon Data Firehose delivery stream. Firehose buffers incoming records and delivers them durably to Amazon S3; it can invoke a Lambda function to transform records before delivery. This removes the single on-premises Kafka-server dependency that caused the data-loss event and avoids administering availability, failover, and capacity for self-managed streaming infrastructure. The service integration also creates a direct ingestion-to-storage path appropriate for high-volume sensor telemetry. AWS documents MQTT support and the managed message broker capabilities in AWS IoT Core message broker. The IoT rule action that writes messages to Firehose is documented in Amazon Data Firehose rule action.

QUESTION NO: 117

A company is configuring connectivity to a multi-account AWS environment to support application workloads fiat serve users in a single geographic region. The workloads depend on a highly available, on-premises legacy system deployed across two locations It is critical for the AWS workloads to manias connectivity to the legacy system, and a minimum of 5 Gbps of bandwidth is required All application workloads within AWS must have connectivity with one another.

Which solution will meet these requirements?

A. Configure multiple AWS Direct Connect (OX) 10 Gbps dedicated connections from a DX partner for each on-premises location Create private virtual interfaces on each connection for each AWS account VPC Associate me private virtual interface with a virtual private gateway attached to each VPC

B. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location Create and attach a virtual private gateway for each AWS account VPC. Create a DX gateway m a central network account and associate it with the virtual private gateways Create a public virtual interface on each DX connection and associate the interface with me DX gateway.

C. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location Create a transit gateway and a DX gateway in a central network account. Create a transit virtual interface for each DX interlace and associate them with the DX gateway. Create a gateway association between the DX gateway and the transit gateway

D. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from a DX partner for each on-premises location Create and attach a virtual private gateway for each AWS account VPC. Create a transit gateway in a central network account and associate It with the virtual private gateways Create a transit virtual interface on each DX connection and attach the interface to the transit gateway.

ANSWER: C

Explanation:

Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location Create a transit gateway and a DX gateway in a central network account. Create a transit virtual interface for each DX interlace and associate them with the DX gateway. Create a gateway association between the DX gateway and the transit gateway is the appropriate architecture because it provides both resilient hybrid connectivity and scalable connectivity among VPCs in multiple AWS accounts. A transit gateway acts as the regional routing hub: application VPCs can be attached to it, including through AWS RAM sharing from a central network account, so that workloads can communicate using centrally managed routing.

A transit virtual interface connects each Direct Connect connection to a Direct Connect gateway. The Direct Connect gateway can then be associated with the transit gateway, extending private on-premises connectivity to the attached VPCs. Using multiple 10 Gbps dedicated connections through two Direct Connect partners and across both on-premises locations supplies capacity well above the 5 Gbps requirement and provides independent connectivity paths for high availability. This follows AWS Direct Connect resiliency guidance, which recommends redundant connections and diverse connectivity locations or providers for critical workloads. See the AWS documentation for Direct Connect gateway associations with transit gateways and the AWS Direct Connect Resiliency Toolkit.

QUESTION NO: 118

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 10 meet these requirements? (Choose 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 accounts 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 Cognoto 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.” establishes a governed multi-account foundation. AWS Control Tower builds on AWS Organizations and supports account provisioning and enrollment, which is appropriate for organizing separate development, staging, production, and shared-network accounts under centralized governance.

“Set up and enable AWS IAM Identity Center (AWS Single Sign-On). Create appropriate permission sets with required MFA for existing accounts.” provides centralized workforce authentication and authorization for AWS accounts and supported applications. IAM Identity Center permission sets can be assigned to groups for account-specific access, while MFA can be required during sign-in through the identity source configuration.

“Create transit gateways and transit gateway VPC attachments in each account. Configure appropriate route tables.” provides private, scalable inter-VPC connectivity. Transit Gateway route tables can segment connectivity: routes can permit production and shared-network VPCs to communicate with all required VPCs, while associating and propagating development and staging attachments so that those environments communicate only with each other. This delivers network isolation through routing policy rather than relying on public connectivity. See AWS Control Tower documentation and Transit Gateway route table documentation.

QUESTION NO: 119

A company needs a hybrid DNS architecture. The architecture must include the company’s on-premises network and a VPC. An AWS Site-to-Site VPN connection connects the VPC to the on-premises network. The company already hosts the onprem.mydc.com domain name on premises.

The company wants to host the myvpc.example.com domain name in the company's AWS account and resolve it to the VPC. The company also needs the on-premises devices to resolve DNS queries to the myvpc.example.com domain.

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

A. Create an Amazon Route 53 private hosted zone for the myvpc.example.com domain. Associate the domain with the VPC.

B. Create an Amazon Route 53 public hosted zone for the myvpc.example.com domain. Associate the domain with the VPC.

C. Use Amazon Route 53 Resolver to create an inbound endpoint in the AWS Region of the VPC.

D. Use Amazon Route 53 Resolver to create an outbound endpoint in the AWS Region of the VPC.

E. Use Amazon Route 53 Resolver to create a forwarding rule for the Route 53 private hosted zone domain and IP addresses of the outbound endpoint.

F. Configure the on-premises DNS resolvers with a conditional forwarding rule for DNS queries for the Amazon Route 53 private hosted zone domain and IP addresses of the inbound endpoint.

ANSWER: A C F

Explanation:

Creating an Amazon Route 53 private hosted zone for `myvpc.example.com` and associating it with the VPC provides an AWS-managed private DNS namespace for records that must resolve within that VPC. Private hosted zone records are available through the Route 53 Resolver associated with the VPC and are not published to the public DNS internet.

Creating a Route 53 Resolver inbound endpoint enables DNS query traffic that originates outside AWS to enter the VPC's Route 53 Resolver service. The endpoint has IP addresses in VPC subnets, which on-premises DNS servers can reach through the Site-to-Site VPN. This is the required direction because the queries originate on premises and need answers from the Route 53 private hosted zone in AWS.

Configuring the on-premises DNS resolvers with a conditional forwarding rule for `myvpc.example.com`, using the inbound endpoint IP addresses as targets, directs only requests for that AWS private namespace across the VPN. Route 53 Resolver then evaluates the associated private hosted zone and returns the appropriate private record response to the on-premises resolver. This design preserves on-premises authority for `onprem.mydc.com` while enabling resolution of the AWS namespace. See Route 53 private hosted zones and Resolver inbound endpoints.

QUESTION NO: 120

A company plans to migrate a legacy on-premises application to AWS. The application is a Java web application that runs on Apache Tomcat with a PostgreSQL database.

The company does not have access to the source code but can deploy the application Java Archive (JAR) files. The application has increased traffic at the end of each month.

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

A. Launch Amazon EC2 instances in multiple Availability Zones. Deploy Tomcat and PostgreSQL to all the instances by using Amazon EFS mount points. Use AWS Step Functions to deploy additional EC2 instances to scale for increased traffic.

B. Provision Amazon EKS in an Auto Scaling group across multiple AWS Regions. Deploy Tomcat and PostgreSQL in the container images. Use a Network Load Balancer to scale for increased traffic.

C. Refactor the Java application into Python-based containers. Use AWS Lambda functions for the application logic. Store application data in Amazon DynamoDB global tables. Use AWS Storage Gateway and Lambda concurrency to scale for increased traffic.

D. Use AWS Elastic Beanstalk to deploy the Tomcat servers with auto scaling in multiple Availability Zones. Store application data in an Amazon RDS for PostgreSQL database. Deploy Amazon CloudFront and an Application Load Balancer to scale for increased traffic.

ANSWER: D

Explanation:

Use AWS Elastic Beanstalk to deploy the Tomcat servers with auto scaling in multiple Availability Zones. Store application data in an Amazon RDS for PostgreSQL database. Deploy Amazon CloudFront and an Application Load Balancer to scale for increased traffic. This approach is the best fit because it retains the application's established Java/Tomcat and PostgreSQL architecture while transferring much of the infrastructure management to AWS managed services. Elastic Beanstalk is designed to deploy and operate web application environments, handling capacity provisioning, Auto Scaling, load-balancer integration, health monitoring, and application deployments. This enables the company to deploy its existing Java application artifacts without source-code access or a substantial architectural rewrite. An Elastic Beanstalk environment can use an Application Load Balancer and instances distributed across multiple Availability Zones, providing scalable and highly available application capacity for predictable month-end demand. Amazon RDS for PostgreSQL removes routine database-administration work, including managed backups, software patching, monitoring, and optional Multi-AZ deployment. CloudFront can cache eligible static and cacheable responses at edge locations, reducing origin load and improving user-facing performance during traffic peaks. Together, these services provide a migration path with limited operational responsibility while preserving the required application runtime and database engine. See AWS Elastic Beanstalk platforms and Amazon RDS for PostgreSQL documentation.

QUESTION NO: 121

A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes.

The company must maintain backups in a separate AWS Region. The company must be able to recover the EC2 instances and their configuration within I business day, with loss of no more than I day ' s worth of data. The company has limited staff and needs a backup solution that optimizes operational efficiency and cost. The company already has created an AWS CloudFormation template that can deploy the required network configuration in a secondary Region.

Which solution will meet these requirements?

A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbooks. Copy the snapshots to the secondary Region. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.

B. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volumes. In the event of a failure, launch theCloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.

C. Use AWS Backup to create a scheduled daily backup plan for the EC2 instances. Configure the backup task to copy the backups to a vault in the secondary Region. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.

D. Deploy EC2 instances of the same size and configuration to the secondary Region. Configure AWS DataSync daily to copy data from the primary Region to the secondary Region. In the event of a failure, launch the CloudFormation template and transfer usage to the secondaryRegion.

ANSWER: C

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. This approach directly meets the one-day recovery point objective by creating daily recovery points and copying them to a backup vault in another Region. AWS Backup is a managed, policy-based service that centrally schedules backups, manages retention, monitors backup jobs, and performs cross-Region copies, minimizing the manual effort required from the company's limited staff.

AWS Backup supports Amazon EC2 instance backups, which protect the instance and its attached EBS volumes. The EC2 restore process can recreate an instance from the recovery point, including relevant instance configuration and volumes, allowing recovery without manually coordinating snapshots across multiple volumes. In a Regional failure, the existing

CloudFormation template first establishes the network infrastructure in the secondary Region. The copied recovery point can then be restored there, enabling the company to meet the one-business-day recovery objective while avoiding the ongoing cost of operating a fully provisioned standby fleet. AWS Backup documentation describes its centralized backup plans and cross-Region copy capabilities, while the EC2 restore guide details recovery of EC2 instances from backup recovery points. See AWS Backup Developer Guide and Restoring Amazon EC2 instances with AWS Backup.

QUESTION NO: 122

A solutions architect is evaluating the reliability of a recently migrated application running on AWS. The front end is hosted on Amazon S3 and accelerated by Amazon CloudFront. The application layer is running in a stateless Docker container on an Amazon EC2 On-Demand Instance with an Elastic IP address. The storage layer is a MongoDB database running on an EC2 Reserved Instance in the same Availability Zone as the application layer.

Which combination of steps should the solutions architect take to eliminate single points of failure with minimal application code changes? (Select TWO.)

A. Create a REST API in Amazon API Gateway and use AWS Lambda functions as the application layer.

B. Create an Application Load Balancer and migrate the Docker container to AWS Fargate.

C. Migrate the storage layer to Amazon DynamoDB.

D. Migrate the storage layer to Amazon DocumentDB (with MongoDB compatibility).

E. Create an Application Load Balancer and move the storage layer to an EC2 Auto Scaling group.

ANSWER: B D

Explanation:

Create an Application Load Balancer and migrate the Docker container to AWS Fargate provides a resilient replacement for the single EC2-hosted application container and its directly associated Elastic IP address. The stateless container can be deployed as an Amazon ECS service using Fargate tasks, with multiple task replicas distributed across Availability Zones. An Application Load Balancer provides one highly available endpoint and routes requests only to healthy tasks. This approach retains the existing Docker deployment model, so it generally requires minimal application code changes while removing dependence on an individual compute instance. AWS documents Fargate as a serverless compute engine for containers and describes using Application Load Balancers with ECS services: AWS Fargate.

Migrate the storage layer to Amazon DocumentDB (with MongoDB compatibility) replaces the single-instance MongoDB deployment with a managed, MongoDB-compatible database service. Amazon DocumentDB cluster storage is replicated across multiple Availability Zones, and replica instances can provide read scaling and automatic failover. Because DocumentDB supports MongoDB APIs and drivers, a MongoDB-oriented application can usually preserve its data-access pattern with limited code modification, subject to compatibility validation. This removes the database host and single-Availability-Zone dependency while providing managed durability and availability. See What is Amazon DocumentDB?.

QUESTION NO: 123

A company ' s solutions architect needs to provide secure Remote Desktop connectivity to users for Amazon EC2 Windows instances that are hosted in a VPC.

The solution must integrate centralized user management with the company ' s on-premises Active Directory. Connectivity to the VPC is through the internet. The company has hardware that can be used to establish an AWS Site-to-Site VPN connection.

Which solution will meet these requirements MOST cost-effectively?

A. Deploy a managed Active Directory by using AWS Directory Service for Microsoft Active Directory. Establish a trust with the on-premises Active Directory.Deploy an EC2 instance as a bastion host in the VP

B. Configure AWS IAM Identity Center (AWS Single Sign-On) to integrate with the on-premises Active Directory by using the AWS Directory Service for MicrosoftActive Directory AD Connector. Configure permission sets against user groups for access to AWS Systems Manager. Use Systems Manager Fleet Manager toaccess the target instances through RDP.

C. Implement a VPN between the on-premises environment and the target VP

D. Deploy a managed Active Directory by using AWS Directory Service for Microsoft Active Directory. Establish a trust with the on-premises Active Directory.Deploy a Remote Desktop Gateway on AWS by using an AWS Quick Start. Ensure that the Remote Desktop Gateway is joined to the domain. Use the Remote Desktop Gateway to access the target instances through RDP.

ANSWER: C

Explanation:

Implement a VPN between the on-premises environment and the target VP is the most cost-effective solution because the company already has compatible hardware for an AWS Site-to-Site VPN connection. AWS Site-to-Site VPN establishes encrypted IPsec tunnels across the internet between the on-premises network and the VPC. This gives corporate users private network paths to the Windows instances, so RDP does not need to be exposed publicly through internet-facing security-group rules or public IP addresses.

Because the VPN provides network connectivity back to the existing on-premises environment, the EC2 Windows instances can join and authenticate against the company's existing Active Directory domain. Administrators can then centrally manage RDP authorization through existing Active Directory users and groups, preserving the company's established identity-management processes without creating another directory or deploying additional remote-access infrastructure. Appropriate VPC route tables, security groups, DNS resolution, and Active Directory ports must be configured so that instances can communicate with the domain controllers over the VPN. AWS documents Site-to-Site VPN as secure connectivity between an on-premises network and a VPC: AWS Site-to-Site VPN User Guide. For domain integration of Windows EC2 instances, see Join an EC2 instance to a Microsoft AD directory.

QUESTION NO: 124

A video processing company uses an AWS Lambda function to handle image processing tasks. An Amazon EventBridge rule that matches the event pattern when a new image is uploaded to an Amazon S3 bucket invokes the Lambda function. The processing task initially operated without errors.

The Lambda function now encounters frequent timeout errors. The Lambda function is configured with the maximum timeout value. A solutions architect must refactor the application's architecture to mitigate invocation failures.

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

A. Build a Docker container image with the application code for deployment. Store the container image in Amazon ECR.

B. Build a Docker container image with the application code for deployment. Store the container image in an S3 bucket with S3 Versioning enabled.

C. Create a new Amazon ECS deployment with the Amazon EC2 launch type. Configure the ECS task definition to use the new Docker container image. Configure the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

D. Create a new Amazon ECS deployment with the Fargate launch type. Configure the ECS task definition to use the new Docker container image. Configure EventBridge to invoke an ECS task by using the ECS task definition.

E. Create a new AWS Step Functions state machine. Configure the state machine to use the new Docker container image. Configure the Lambda function to invoke the state machine when a new file arrives in Amazon S3.

ANSWER: A D

Explanation:

"Build a Docker container image with the application code for deployment. Store the container image in Amazon ECR." and "Create a new Amazon ECS deployment with the Fargate launch type. Configure the ECS task definition to use the new Docker container image. Configure EventBridge to invoke an ECS task by using the ECS task definition." provide a managed, event-driven solution for processing workloads that can exceed AWS Lambda's maximum execution duration of 15 minutes. Amazon ECR is AWS's managed container registry and integrates directly with Amazon ECS task definitions, allowing the task to pull the versioned image at runtime. Fargate runs ECS tasks without requiring the company to provision,

patch, scale, or maintain Amazon EC2 container hosts, minimizing operational work. EventBridge rules can use Amazon ECS tasks as targets and can specify Fargate task networking and launch settings through ECS parameters. The existing S3-to-EventBridge event pattern can therefore trigger an independently running Fargate task directly, removing Lambda from the long-running processing path. This supports image-processing jobs whose duration is not constrained by Lambda's timeout while retaining managed AWS services and decoupled event-based invocation. See the AWS Lambda timeout documentation and Amazon EventBridge target documentation.

QUESTION NO: 125

A solutions architect must update an application environment within AWS Elastic Beanstalk using a With green deployment methodology. The solutions architect creates an environment that is identical to the existing application environment and deploys the application to the new environment.

What should be done next to complete the update?

A. Redirect to the new environment using Amazon Route 53

B. Select the Swap Environment URLs option.

C. Replace the Auto Scaling launch configuration

D. Update the DNS records to point to the green environment

ANSWER: B

Explanation:

"Select the Swap Environment URLs option." is the appropriate next step for an AWS Elastic Beanstalk blue/green deployment. The existing environment continues to serve the production application while the separately created environment runs the updated application version. Before the cutover, the new environment should be validated through its own environment URL to confirm that application behavior, configuration, scaling, and dependencies operate as intended.

After validation, Elastic Beanstalk's environment URL swap exchanges the CNAMEs assigned to the two environments. Consequently, requests sent to the original production Elastic Beanstalk environment URL are directed to the new, updated environment without requiring clients to change the URL they use. The former production environment remains available with the previous version, which provides a rapid rollback mechanism: the environment URLs can be swapped again if an issue is discovered after cutover. AWS documents CNAME swapping as the core traffic-switching action for Elastic Beanstalk blue/green deployments. For details, see Elastic Beanstalk blue/green deployments and CNAME swapping and deploying an existing application version.

QUESTION NO: 126

A company hosts an application on AWS. The application uses AWS Lambda functions that are invoked by an Amazon API Gateway API. The company has an Amazon CloudFront distribution that uses the API Gateway API as its origin. The CloudFront distribution serves web requests to customers worldwide.

During testing, users experienced slow responses from the application APIs. The company discovered that requests from different AWS Regions contained inconsistent query parameters with mixed-case letters, which caused increased cache misses and more requests to reach the Lambda functions.

The company wants to ensure that the API consistently provides responses with minimal latency.

Which solution will meet these requirements?

A. Create a new Lambda function to sort incoming request query parameters alphabetically and convert the parameters to lowercase. Configure the CloudFront distribution to use the Lambda@Edge function type. Configure the Lambda function to invoke on origin request.

B. Create a CloudFront function to sort incoming request query parameters alphabetically and convert the parameters to lowercase. Configure the CloudFront distribution to use the CloudFront Functions function type. Configure the CloudFront function to invoke on viewer request.

C. Configure the API Gateway API to use mapping templates to sort incoming request query parameters alphabetically and convert the parameters to lowercase before Lambda processes the request.

D. Configure the API Gateway API to use a Lambda authorizer to sort incoming request query parameters alphabetically and convert the parameters to lowercase before Lambda processes the request.

ANSWER: B

Explanation:

Create a CloudFront function to sort incoming request query parameters alphabetically and convert the parameters to lowercase. Configure the CloudFront distribution to use the CloudFront Functions function type. Configure the CloudFront function to invoke on viewer request. A viewer-request CloudFront Function executes at the edge before CloudFront performs its cache lookup. It can normalize the query string by applying a consistent lowercase representation and parameter order, so requests that are logically equivalent generate the same cache key. This increases the cache-hit ratio and prevents unnecessary forwarding to the API Gateway origin and Lambda functions.

CloudFront Functions are intended for lightweight JavaScript transformations on viewer-request and viewer-response events, including URL and query-string manipulation. Their edge execution model is particularly appropriate when the goal is to minimize latency globally: cached responses can be served directly from the edge location after the normalized request is evaluated. This approach also centralizes normalization in CloudFront, ensuring that requests from all Regions are treated consistently before any origin processing occurs. The cache behavior must include the relevant query strings in the cache key so that CloudFront can distinguish requests whose normalized parameter values differ. See the AWS documentation for CloudFront Functions and CloudFront cache keys.

QUESTION NO: 127

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? (Select 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:

Deploying the security tool on EC2 instances in a new Auto Scaling group in the existing VPC enables the company to retain its established third-party packet-inspection software while operating it as a fleet of virtual network appliances. The Auto Scaling group can maintain capacity and replace failed appliance instances. Deploying appliance instances across multiple Availability Zones provides Regional high availability and allows inspection capacity to scale independently from the web application.

Provisioning a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool provides the purpose-built service insertion mechanism for these virtual appliances. Gateway Load Balancer distributes traffic flows to security appliance targets transparently at the network layer, using GENEVE encapsulation, and preserves flow symmetry so

that packets belonging to a connection are consistently processed by the appropriate appliance. Gateway Load Balancer endpoints and VPC route tables can direct ingress and egress paths through the inspection fleet, supporting real-time inspection without requiring changes to the application instances. Using zonal endpoints and targets in multiple Availability Zones creates a resilient inspection architecture while keeping traffic local to an Availability Zone where possible. AWS documents Gateway Load Balancer specifically for deploying, scaling, and managing third-party virtual appliances such as firewalls, intrusion detection systems, and deep packet inspection systems. See Gateway Load Balancer and Gateway Load Balancer endpoints.

QUESTION NO: 128

A company manages an on-premises JavaScript front-end web application. The application is hosted on two servers secured with a corporate Active Directory. The application calls a set of Java-based microservices on an application server and stores data in a clustered MySQL database. The application is heavily used during the day on weekdays. It is lightly used during the evenings and weekends.

Daytime traffic to the application has increased rapidly, and reliability has diminished as a result. The company wants to migrate the application to AWS with a solution that eliminates the need for server maintenance, with an API to securely connect to the microservices.

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

A. Host the web application on Amazon S3. Use Amazon Cognito identity pools (federated identities) with SAML for authentication and authorization.

B. Host the web application on Amazon EC2 with Auto Scaling. Use Amazon Cognito federation and Login with Amazon for authentication and authorization.

C. Create an API layer with Amazon API Gateway. Rehost the microservices on AWS Fargate containers.

D. Create an API layer with Amazon API Gateway. Rehost the microservices on Amazon Elastic Container Service (Amazon ECS) containers.

E. Replatform the database to Amazon RDS for MySQL.

F. Replatform the database to Amazon Aurora MySQL Serverless.

ANSWER: A C F

Explanation:

“Host the web application on Amazon S3. Use Amazon Cognito identity pools (federated identities) with SAML for authentication and authorization.” provides a serverless hosting model for the JavaScript front end. Amazon S3 can serve static web assets without web-server administration, and Cognito federation with SAML allows the existing corporate identity system to participate in authentication and authorization.

“Create an API layer with Amazon API Gateway. Rehost the microservices on AWS Fargate containers.” supplies a managed and secure API entry point while moving the Java microservices into containers that do not require the company to provision or maintain container-host EC2 instances. API Gateway can expose and secure application APIs, and Fargate runs the containerized workloads with server infrastructure managed by AWS. AWS describes this operating model in its AWS Fargate documentation.

“Replatform the database to Amazon Aurora MySQL Serverless.” retains MySQL compatibility while using managed, automatically scaling database capacity. This aligns with substantial weekday demand and much lower off-hours demand, reducing database capacity-management effort while improving responsiveness to variable traffic. Aurora Serverless is documented by AWS at Amazon Aurora Serverless.

QUESTION NO: 129

A company has five development teams that have each created five AWS accounts to develop and host applications. To track spending, the development teams log in to each account every month, record the current cost from the AWS Billing and Cost Management console, and provide the information to the company's finance team. The company has strict compliance

requirements and needs to ensure that resources are created only in AWS Regions in the United States. However, some resources have been created in other Regions. A solutions architect needs to implement a solution that gives the finance team the ability to track and consolidate expenditures for all the accounts. The solution also must ensure that the company can create resources only in Regions in the United States. Which combination of steps will meet these requirements in the MOST operationally efficient way? (Choose three.)

A. Create a new account to serve as a management account. Create an Amazon S3 bucket for the finance team. Use AWS Cost and Usage Reports to create monthly reports and to store the data in the finance team's S3 bucket.

B. Create a new account to serve as a management account. Deploy an organization in AWS Organizations with all features enabled. Invite all the existing accounts to the organization. Ensure that each account accepts the invitation.

C. Create an OU that includes all the development teams. Create an SCP that allows the creation of resources only in Regions that are in the United States. Apply the SCP to the OU.

D. Create an OU that includes all the development teams. Create an SCP that denies the creation of resources in Regions that are outside the United States. Apply the SCP to the OU.

E. Create an IAM role in the management account. Attach a policy that includes permissions to view the Billing and Cost Management console. Allow the finance team users to assume the role. Use AWS Cost Explorer and the Billing and Cost Management console to analyze cost.

F. Create an IAM role in each AWS account. Attach a policy that includes permissions to view the Billing and Cost Management console. Allow the finance team users to assume the role.

ANSWER: B D E

Explanation:

Creating a new account to serve as a management account, deploying an AWS Organization with all features enabled, and inviting the existing accounts establishes centralized governance and consolidated billing. With consolidated billing, charges from member accounts are combined under the management account, eliminating the monthly manual collection of cost information from every development account. AWS Organizations must use all features to support service control policies (SCPs).

Creating an OU for the development accounts and applying an SCP that denies resource creation outside United States Regions provides the required preventive control at scale. An explicit deny, commonly implemented with the `aws:RequestedRegion` global condition key and an approved-Region list, overrides otherwise-permitted IAM actions. Applying that guardrail to the OU ensures it is inherited by all development accounts placed in that OU.

Creating an IAM role in the management account for finance users centralizes access to organization-wide billing data. With the appropriate Billing and Cost Management and Cost Explorer permissions, the finance team can assume that role and analyze consolidated organizational costs from one account. This combination minimizes repetitive account-by-account administration while applying consistent regional compliance controls. See AWS Organizations and consolidated billing and SCP examples for restricting AWS Regions.

QUESTION NO: 130

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' 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 and configure the app to use its acceleration endpoint for uploads. S3 Transfer Acceleration is designed for transferring data over long geographic distances. It accepts uploads at an AWS edge location close to the client and moves the data over the AWS global network to the destination bucket in us-east-1. This can improve upload performance for users in Australia who are sending files to a bucket in the United States. The application must use the bucket's Transfer Acceleration endpoint for the feature to apply.

Configure the app to break video files into chunks and use multipart upload to transfer files to Amazon S3. Multipart upload is appropriate for these large video objects because the app can upload parts independently and in parallel. If a network interruption causes an individual part to fail, the app retries only that part rather than restarting the entire 1 GB to 10 GB upload. This improves both transfer efficiency and resiliency for unreliable or long-distance mobile connections. Amazon S3 requires multipart upload for objects larger than 5 GB and recommends it for large-object uploads. See Amazon S3 Transfer Acceleration and Amazon S3 multipart upload overview.

QUESTION NO: 131

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:

Cross-account reads of Amazon S3 objects encrypted with a customer managed AWS KMS key require authorization from both Amazon S3 and AWS KMS. "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" establishes the S3 resource-policy authorization for the

external account. The Strategy account can then delegate that granted access to its IAM principals, including the assumed `strategy_reviewer` role.

“Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key” applies least privilege to the role: it needs object-read actions such as `s3:GetObject` for the shared content and `kms:Decrypt` to consume objects protected by SSE-KMS. “Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role” is also necessary because key policies are the primary authorization mechanism for a KMS key. It explicitly permits the cross-account role to use the Creative account’s key for decryption. Together, these permissions allow only object viewing and required cryptographic decryption. See Amazon S3 cross-account access guidance and AWS KMS cross-account key access guidance.

QUESTION NO: 132

A company exposes a public REST API through an Amazon API Gateway Regional REST API. The company must require clients to present certificates that are issued by the company before clients can establish a connection to the API. The company does not want to implement certificate validation in Lambda functions or in the backend application.

Which solution will meet these requirements?

A. Configure a custom domain name for the Regional REST API. Store the certificate authority truststore in Amazon S3. Configure mutual TLS authentication on the custom domain name.

B. Create an AWS WAF web ACL with an IP set rule. Associate the web ACL with the API Gateway API.

C. Create an API Gateway usage plan and require an API key for every API method.

D. Configure AWS Lambda authorizers to validate each client certificate before invoking the API backend.

ANSWER: A

Explanation:

Configure a custom domain name for the Regional REST API. Create an Amazon S3 bucket that contains a truststore with the company certificate authority certificates. Configure mutual TLS authentication on the custom domain name and specify the S3 truststore location.

API Gateway mutual TLS authentication validates client certificates during the TLS handshake before API Gateway processes the request or invokes an integration. The truststore contains the trusted certificate authority certificates that API Gateway uses to validate the client certificate chain. This provides certificate-based client authentication at the API endpoint and avoids custom code in Lambda functions or backend services. The custom domain must use a Regional endpoint and an ACM certificate for the server side of the TLS connection. See Mutual TLS authentication for REST APIs.

QUESTION NO: 133

A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.

Which solution will meet these requirements while providing the FASTEST storage performance?

A. Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluster. Configure the ReplicaSet to mount the file system. Direct the application to store files in the file system. Configure AWS Backup to back up and retain copies of the data for 1 year.

B. Create an Amazon Elastic Block Store (Amazon EBS) volume. Enable the EBS Multi-Attach feature. Configure the ReplicaSet to mount the EBS volume. Direct the application to store files inthe EBS volume. Configure AWS Backup to back up and retain copies of the data for 1 year.

C. Create an Amazon S3 bucket. Configure the ReplicaSet to mount the S3 bucket. Direct the application to store files in the S3 bucket. Configure S3 Versioning to retain copies of the data. Configure an S3 Lifecycle policy to delete objects after 1 year.

D. Configure the ReplicaSet to use the storage available on each of the running application pods to store the files locally. Use a third-party tool to back up the EKS cluster for 1 year.

ANSWER: A

Explanation:

Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluster. Configure the ReplicaSet to mount the file system. Direct the application to store files in the file system. Configure AWS Backup to back up and retain copies of the data for 1 year. This solution provides the required shared, POSIX-compatible file storage, allowing all application pods to concurrently read and write the same files even when the pods run on nodes in different Availability Zones. EFS is a Regional service and is designed for concurrent access from multiple compute instances; for optimal connectivity and latency, mount targets should be available in the Availability Zones used by the EKS worker nodes. The Amazon EFS CSI driver provides the Kubernetes integration needed to expose EFS as persistent storage to pods in EKS. EFS is especially appropriate for a workload that creates many small files and requires a conventional shared file-system interface rather than object storage semantics. AWS Backup natively supports EFS and can apply a backup plan with a one-year retention period, providing managed, policy-based protection without requiring custom backup software. AWS documents using EFS volumes with EKS through the Amazon EFS CSI driver and describes EFS backup management in the AWS Backup Developer Guide.

QUESTION NO: 134

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. improves transfer performance for clients that are geographically distant from the destination bucket. Instead of sending upload traffic entirely across the public internet to us-east-1, users in Australia can connect to a nearby AWS edge location. AWS then carries the data over its optimized global network to the target S3 bucket. This is designed for fast, long-distance transfers to a centralized S3 bucket and requires the application to use the bucket's acceleration endpoint. See the Amazon S3 Transfer Acceleration documentation.

Configure the app to break the video files into chunks. Use a multipart upload to transfer files to Amazon S3. increases reliability and can improve throughput for objects in the 1 GB to 10 GB range. Multipart upload sends an object as independently uploaded parts, allowing the client to upload multiple parts in parallel when appropriate. Most importantly for users with variable or unreliable long-distance connectivity, a failed transfer requires retrying only the affected part rather than restarting the entire large-file upload. After all parts are successfully uploaded, S3 assembles them into the final object. AWS recommends multipart upload for large objects and for improving recovery from network failures. See Uploading and copying objects using multipart upload.

QUESTION NO: 135

A company wants to create a single Amazon S3 bucket for its data scientists to store work-related documents. The company uses AWS IAM Identity Center to authenticate all users. A group for the data scientists was created. The company wants to give the data scientists access to only their own work. The company also wants to create monthly reports that show which documents each user accessed. Which combination of steps will meet these requirements? (Choose two.)

A. Create a custom IAM Identity Center permission set to grant the data scientists access to an S3 bucket prefix that matches their username tag. Use a policy to limit access to paths with the ${aws:PrincipalTag/userName}/* condition.

B. Create an IAM Identity Center role for the data scientists group that has Amazon S3 read access and write access. Add an S3 bucket policy that allows access to the IAM Identity Center role.

C. Configure AWS CloudTrail to log S3 data events and deliver the logs to an S3 bucket. Use Amazon Athena to run queries on the CloudTrail logs in Amazon S3 and generate reports.

D. Configure AWS CloudTrail to log S3 management events to CloudWatch. Use Amazon Athena's CloudWatch connector to query the logs and generate reports.

E. Enable S3 access logging to EMR File System (EMRFS). Use Amazon S3 Select to query logs and generate reports.

ANSWER: A C

Explanation:

Create a custom IAM Identity Center permission set to grant the data scientists access to an S3 bucket prefix that matches their username tag. Use a policy to limit access to paths with the ${aws:PrincipalTag/userName}/* condition. This uses attribute-based access control (ABAC) to provide per-user isolation while retaining one shared bucket. IAM Identity Center can pass configured user attributes as principal tags in the role session created by a permission set. An IAM policy can then use the principal-tag policy variable in S3 object ARNs and in an S3 prefix condition for listing, so each authenticated user can work only with objects under that user's designated prefix. This avoids creating separate buckets or static policies for each data scientist.

Configure AWS CloudTrail to log S3 data events and deliver the logs to an S3 bucket. Use Amazon Athena to run queries on the CloudTrail logs in Amazon S3 and generate reports. CloudTrail data events record object-level S3 API activity, including object reads and writes, together with the requesting identity, object key, event time, and source details. Delivering the trail to S3 creates durable log files that Athena can query using SQL. Scheduled Athena queries or report-generation processes can therefore produce monthly, per-user reports of documents accessed. AWS documents IAM Identity Center ABAC at IAM Identity Center ABAC and S3 data-event logging at CloudTrail data events.

QUESTION NO: 136

A company wants to migrate an Amazon Aurora MySQL DB cluster from an existing AWS account to a new AWS account in the same AWS Region. Both accounts are members of the same organization in AWS Organizations. The company must minimize database service interruption before the company performs DNS cutover to the new database. Which migration strategy will meet this requirement? (Choose two.)

A. Take a snapshot of the existing Aurora database. Share the snapshot with the new AWS account. Create an Aurora DB cluster in the new account from the snapshot.

B. Create an Aurora DB cluster in the new AWS account. Use AWS Database Migration Service (AWS DMS) to migrate data between the two Aurora DB clusters.

C. Use AWS Backup to share an Aurora database backup from the existing AWS account to the new AWS account. Create an Aurora DB cluster in the new AWS account from the snapshot.

D. Create an Aurora DB cluster in the new AWS account. Use AWS Application Migration Service to migrate data between the two Aurora DB clusters.

ANSWER: A B

Explanation:

Taking a snapshot of the existing Aurora database, sharing the snapshot with the new AWS account, and creating an Aurora DB cluster in the new account from the snapshot is a supported cross-account Aurora migration approach. The source cluster can continue serving application traffic while the snapshot is shared and the destination cluster is provisioned. This enables the company to complete the bulk of migration preparation before DNS cutover. For encrypted snapshots, the source account must use a customer managed AWS KMS key and allow the destination account to use that key as part of sharing the snapshot. Aurora documents the required snapshot-sharing process at Sharing DB cluster snapshots.

Creating an Aurora DB cluster in the new AWS account and using AWS Database Migration Service (AWS DMS) to migrate data between the two Aurora DB clusters also minimizes interruption. AWS DMS can perform an initial full load and then use change data capture to replicate ongoing source changes to the target. The target can therefore be prepared and kept close to current while the original database remains in service; the final cutover is limited to the time needed to stop or reconcile writes and switch DNS. AWS documents this full-load-and-CDC migration pattern at AWS DMS change data capture.

QUESTION NO: 137

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 win deploy a two-net application In a new VPC To simplify the configuration, the db.example com CNAME record set tor 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 /eto/resolv.conf file

C. Create an authorization lo 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 m Account B Configure Route 53 replication between AWS accounts

E. Associate a new VPC in Account B with a hosted zone in Account

ANSWER: C E

Explanation:

A Route 53 private hosted zone resolves records only for VPCs that are associated with that hosted zone. Because the private hosted zone is owned by Account A while the application instance is in a new VPC in Account B, Account A must first explicitly authorize the cross-account VPC association. “Create an authorization lo associate the private hosted zone in Account A with the new VPC In Account B” provides this required authorization. The authorization identifies the target VPC and permits the hosted-zone owner to establish the association across accounts.

After the authorization exists, “Associate a new VPC in Account B with a hosted zone in Account” performs the VPC association. Once associated, DNS queries from the EC2 instance in that VPC can use the private hosted zone and resolve `db.example.com` to the CNAME record for the Amazon RDS endpoint. The VPC must also have DNS support and DNS hostnames enabled, which are normally enabled by default for a VPC. AWS documents this as a two-step process for associating a Route 53 private hosted zone with a VPC in another AWS account: create the authorization in the hosted-zone account, then create the association from the VPC account. See Associating a private hosted zone with VPCs in different AWS accounts and Route 53 private hosted zones.

QUESTION NO: 138

A company wants to use Amazon Workspaces in combination with thin client devices to replace aging desktops. Employees use the desktops to access applications that work with clinical trial data. Corporate security policy states that access to the applications must be restricted to only company branch office locations. The company is considering adding an additional branch office in the next 6 months.

Which solution meets these requirements with the MOST operational efficiency?

A. Create an IP access control group rule with the list of public addresses from the branch offices. Associate the IP access control group with the Workspaces directory.

B. Use AWS Firewall Manager to create a web ACL rule with an IPSet with the list to public addresses from the branch office Locations-Associate the web ACL with the Workspaces directory.

C. Use AWS Certificate Manager (ACM) to issue trusted device certificates to the machines deployed in the branch office locations. Enable restricted access on the Workspaces directory.

D. Create a custom Workspace image with Windows Firewall configured to restrict access to the public addresses of the branch offices. Use the image to deploy the Workspaces.

ANSWER: A

Explanation:

Create an IP access control group rule with the list of public addresses from the branch offices. Associate the IP access control group with the Workspaces directory. Amazon WorkSpaces IP access control groups are designed specifically to limit WorkSpaces client connections by source IP address or CIDR range. By entering each branch office's stable public egress IP range and associating the group with the relevant WorkSpaces directory, the company ensures that users can connect only while their client traffic originates from an approved office network. This enforcement occurs at the WorkSpaces service level before the desktop session is established, which is appropriate for protecting access to applications handling clinical trial data. It also avoids deploying and maintaining endpoint-specific network controls on every thin client or every virtual desktop. When the new branch office opens, an administrator can add that location's public egress IP address or CIDR range to the existing IP access control group, without rebuilding WorkSpaces or changing the desktop image. AWS documents that IP access control groups specify trusted IP address ranges and can be associated with WorkSpaces directories to control where users connect from. See Amazon WorkSpaces IP access control groups and WorkSpaces directory configuration.

QUESTION NO: 139

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 ' ' (Select 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 ' sVPC

ANSWER: A C

Explanation:

Create a transit gateway Attach all the company ' s VPCs and relevant subnets to the transit gateway provides a centrally managed, scalable hub for connectivity among the company's VPCs, including VPCs in different AWS accounts and Regions. Transit Gateway avoids the operational burden of maintaining a growing mesh of individual connections. Each VPC attachment can use route tables to define the intended connectivity, and new VPCs can be incorporated by creating an attachment and updating centralized routing as needed. AWS Transit Gateway is specifically designed to simplify network topology as the number of connected VPCs grows.

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 exposes only the license-validation service privately to customer VPCs. Customers access the service through interface VPC endpoints, while the management VPC does not receive general network-level access back into customer VPCs. This supports the required one-way service access model and keeps the management resource private, without requiring public IP addressing, internet gateways, or broad VPC-to-VPC routing for that service.

See AWS Transit Gateway documentation and AWS PrivateLink concepts.

QUESTION NO: 140

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

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

ANSWER: C

Explanation:

"Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled." provides scaling for both tiers while preserving session continuity for the stateful application. Aurora Auto Scaling dynamically adjusts the number of Aurora Replicas according to a defined target metric, such as average CPU utilization or active database connections. This increases read capacity as demand grows, allowing application read workloads to be distributed across the Aurora Replica fleet. Aurora's writer instance continues to handle write operations, while replicas provide scalable read capacity. AWS documents this capability in its Aurora Auto Scaling documentation.

An Application Load Balancer is designed for HTTP and HTTPS application traffic and integrates directly with an EC2 Auto Scaling group. Round robin routing distributes new requests across healthy registered targets. Enabling target-group stickiness establishes session affinity, so requests from a user continue to reach the same application instance for the configured stickiness duration. This is important because application session state is held locally rather than externalized to a shared session store. AWS describes ALB routing algorithms and sticky sessions in the Application Load Balancer target groups documentation.

QUESTION NO: 141

A company has an asynchronous HTTP application that is hosted as an AWS Lambda function. A public Amazon API Gateway endpoint invokes the Lambda function. The Lambda function and the API Gateway endpoint reside in the us-east-1 Region. A solutions architect needs to redesign the application to support failover to another AWS Region. Which solution will meet these requirements?

A. Create an API Gateway endpoint in the us-west-2 Region to direct traffic to the Lambda function in us-east-1. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.

B. Create an Amazon Simple Queue Service (Amazon SQS) queue. Configure API Gateway to direct traffic to the SQS queue instead of to the Lambda function. Configure the Lambda function to pull messages from the queue for processing.

C. Deploy the Lambda function to the us-west-2 Region. Create an API Gateway endpoint in us-west-2 10 direct traffic to the Lambda function in us-west-2. Configure AWS Global Accelerator and an Application Load Balancer to manage traffic across the two API Gateway endpoints.

D. Deploy the Lambda function and an API Gateway endpoint to the us-west-2 Region. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.

ANSWER: D

Explanation:

Deploy the Lambda function and an API Gateway endpoint to the us-west-2 Region. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints. This creates a complete, independently operating application stack in a second AWS Region: the secondary API Gateway endpoint invokes a Lambda function that is also local to that Region. Consequently, an outage or impairment affecting the primary Region does not leave the failover endpoint dependent on resources in that same impaired Region.

Amazon Route 53 failover routing provides primary and secondary DNS records and evaluates the health of the primary endpoint by using an associated health check. While the primary endpoint is healthy, Route 53 answers DNS queries with the primary record. If the health check determines that the primary is unhealthy, Route 53 returns the secondary record, directing new client DNS resolutions to the API deployed in us-west-2. In practice, the API endpoints should be presented through Regional custom domain names and Route 53 alias records, allowing clients to use one stable DNS name while Route 53 selects the active Regional target. This is an AWS-supported multi-Region active/passive failover pattern. See Amazon Route 53 failover routing and API Gateway Regional custom domain names.

QUESTION NO: 142

A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.

The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.

Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?

A. Deploy a Lambda@Edge function to sort parameters by name and force them to be lowercase. Select the CloudFront viewer request trigger to invoke the function.

B. Update the CloudFront distribution to disable caching based on query string parameters.

C. Deploy a reverse proxy after the load balancer to post-process the emitted URLs in the application to force the URL strings to be lowercase.

D. Update the CloudFront distribution to specify casing-insensitive query string processing.

ANSWER: A

Explanation:

Deploy a Lambda@Edge function to sort parameters by name and force them to be lowercase. Select the CloudFront viewer request trigger to invoke the function. is the appropriate solution because CloudFront evaluates the

viewer request before performing its cache lookup. When a cache policy includes query strings in the cache key, query-string parameter names, values, and ordering can cause requests that are logically equivalent to be treated as separate cacheable objects. Normalizing the query string at the viewer-request event ensures that requests reach cache-key evaluation with one consistent representation.

The function can parse the incoming query string, apply the application's agreed case normalization, and sort parameters deterministically by name before CloudFront attempts to serve the request from cache. This consolidates requests such as the same parameters supplied in different orders into the same cache key, allowing existing edge caching to be used more effectively without waiting for changes to the application's URL-generation behavior. Lambda@Edge is designed to modify viewer requests at CloudFront edge locations, making it a direct way to address cache-key fragmentation globally. AWS documents a Lambda@Edge example specifically for normalizing query-string parameters and explains that CloudFront can cache different query-string combinations separately. See Lambda@Edge query string normalization examples and CloudFront caching based on query string parameters.

QUESTION NO: 143

A company migrated to AWS and uses AWS Business Support. The company wants to monitor the cost-effectiveness of Amazon EC2 instances across AWS accounts. The EC2 instances have tags for department, business unit, and environment. Development EC2 instances have high cost but low utilization. The company needs to detect and stop any underutilized development EC2 instances. Instances are underutilized if they had 10% or less average daily CPU utilization and 5 MB or less network I/O for at least 4 of the past 14 days. Which solution will meet these requirements with the LEAST operational overhead?

A. Configure Amazon CloudWatch dashboards to monitor EC2 instance utilization based on tags for department, business unit, and environment. Create an Amazon EventBridge rule that invokes an AWS Lambda function to stop underutilized development EC2 instances.

B. Configure AWS Systems Manager to track EC2 instance utilization and report underutilized instances to Amazon CloudWatch. Filter the CloudWatch data by tags for department, business unit, and environment. Create an Amazon EventBridge rule that invokes an AWS Lambda function to stop underutilized development EC2 instances.

C. Create an Amazon EventBridge rule to detect low utilization of EC2 instances reported by AWS Trusted Advisor. Configure the rule to invoke an AWS Lambda function that filters the data by tags for department, business unit, and environment and stops underutilized development EC2 instances.

D. Create an AWS Lambda function to run daily to retrieve utilization data for all EC2 instances. Save the data to an Amazon DynamoDB table. Create an Amazon QuickSight dashboard that uses the DynamoDB table as a data source to identify and stop underutilized development EC2 instances.

ANSWER: C

Explanation:

Create an Amazon EventBridge rule to detect low utilization of EC2 instances reported by AWS Trusted Advisor. Configure the rule to invoke an AWS Lambda function that filters the data by tags for department, business unit, and environment and stops underutilized development EC2 instances. This is the lowest-overhead solution because the AWS Trusted Advisor Low Utilization Amazon EC2 Instances cost-optimization check already evaluates the exact utilization pattern in the requirement: average daily CPU utilization of 10% or less and network I/O of 5 MB or less for at least 4 days during the previous 14 days. AWS Business Support provides access to Trusted Advisor cost-optimization checks, so the company does not need to build a separate process to collect Amazon CloudWatch metrics, retain 14-day utilization history, or calculate thresholds itself.

Amazon EventBridge receives Trusted Advisor check-result events and can invoke Lambda automatically when relevant findings are reported. The Lambda function can retrieve the affected EC2 instance's tags, apply the organization's department, business-unit, and environment policy, and stop only development instances. This combines AWS-managed utilization analysis with a small amount of targeted automation, while allowing the same event-driven pattern to be deployed or centrally coordinated across accounts. See AWS Trusted Advisor cost optimization checks and Monitoring Trusted Advisor check results with Amazon EventBridge.

QUESTION NO: 144

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. Addproduction 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:

Creating an AWS Control Tower landing zone in the company's management account provides the required multi-account governance foundation, including a governed landing zone, account baselines, controls, and organizational units. Separating production and development accounts into distinct production and development OUs establishes the boundary needed to apply encryption governance only to production workloads. This OU-based design also ensures that newly enrolled or created production accounts inherit the applicable production governance scope without affecting development accounts.

Existing AWS accounts must be brought under the AWS Organizations organization so they can be enrolled and governed as part of the Control Tower environment. Service control policies can provide preventive governance for encryption requirements by restricting actions that would create noncompliant EBS resources. Applying an EBS-encryption detection guardrail to the production OU adds continuous compliance evaluation for the accounts in that OU. Because Control Tower controls are enabled at the OU level, the detection control applies consistently to current production accounts and accounts added to that OU in the future. AWS documents Control Tower's landing-zone and control model at What is AWS Control Tower? and explains how controls are applied to OUs at AWS Control Tower controls.

QUESTION NO: 145

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 ElastiCache for Memcached in front of the database

B. Use Amazon ElastiCache 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:

Use RDS Proxy in front of the database, migrate the database to Amazon Aurora MySQL, and create an Amazon Aurora Replica. Aurora provides a distributed storage volume that replicates data across multiple Availability Zones and supports highly available database clusters. An Aurora Replica is an already-running database instance that can be promoted to become the writer during a failover, avoiding the need to create or fully initialize a replacement database instance when an outage occurs. This architecture is designed to provide faster recovery for write workloads than a traditional RDS for MySQL Multi-AZ deployment.

RDS Proxy complements Aurora failover by presenting applications with a stable proxy endpoint and managing pooled database connections. During a writer failover, the proxy can preserve and efficiently reestablish connections to the newly promoted writer, reducing application disruption caused by connection handling, DNS changes, and reconnection storms. Together, Aurora MySQL, an Aurora Replica, and RDS Proxy provide both a fast database failover target and reduced client-side connection interruption, supporting the requirement for an outage of less than 20 seconds. See Amazon Aurora high availability and Amazon RDS Proxy documentation.

QUESTION NO: 146

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? (Choose 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. supplies the required dedicated, private network path from the company's on-premises environment to a VPC. A private virtual interface (private VIF) is used to reach resources in a VPC through AWS Direct Connect, rather than sending traffic over the public internet.

Create an Amazon S3 interface endpoint in the networking account. exposes Amazon S3 through private IP addresses in that VPC by using AWS PrivateLink. The on-premises systems can route to those endpoint network interfaces over the Direct Connect private VIF, keeping the S3 data path private. S3 interface endpoints are specifically suitable when S3 must be accessed privately from an on-premises network through Direct Connect or a VPN connection. Access to buckets distributed across the three AWS accounts remains governed by IAM authorization, S3 bucket policies, and the VPC endpoint policy, so the centralized networking VPC can provide connectivity without requiring all buckets to reside in that account. The design can also use appropriate private DNS and on-premises DNS resolution so S3 requests resolve to the interface endpoint addresses. See Amazon S3 interface endpoints and AWS Direct Connect virtual interfaces.

QUESTION NO: 147

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 E F

Explanation:

Creating a new AWS Control Tower landing zone in the company's management account provides the required managed multi-account foundation, including AWS Organizations integration, account governance, built-in blueprints, and Control Tower controls (formerly called guardrails). Separating the accounts into production and development organizational units allows governance to be applied specifically to the production environment. Existing production accounts can be enrolled and placed in the production organizational unit, while future production accounts receive the same governance when they are provisioned or moved into that organizational unit.

Control Tower controls are centrally administered from the management account. Creating the EBS-encryption detection guardrail there establishes the control configuration through the service's supported governance workflow. Applying the guardrail to the production organizational unit scopes the control to every account in that organizational unit, including accounts added later, while leaving the development organizational unit unaffected. The EBS encryption control uses AWS Config-based evaluation to continually identify EBS volumes that do not meet the encryption requirement, providing centralized visibility and compliance governance for the production estate. AWS documents Control Tower controls and their organizational-unit targeting in the AWS Control Tower controls guide and describes landing-zone governance in the How AWS Control Tower works guide.

QUESTION NO: 148

Question:

A company runs a Linux app on Amazon EKS usingM6iEC2 instances under a Savings Plan that is about to expire. They want toreduce costsafter expiration.

A. Rebuild containers forARM64architecture.

B. Rebuild containers for container compatibility (invalid/unclear).

C. Migrate EKS nodes toGraviton(e.g., C7g, M7g).

D. Replace nodes with latestx86_64instances.

E. Purchase a new Savings Plan for the Graviton workload.

F. Purchase new Savings Plan for x86_64 instances.

ANSWER: A C E

Explanation:

Rebuilding containers for ARM64 architecture, migrating EKS nodes to Graviton instances such as C7g or M7g, and purchasing a new Savings Plan for the Graviton workload are an effective cost-optimization path. M6i instances use x86-based Intel processors, whereas C7g and M7g instances are powered by AWS Graviton processors and use the ARM64 instruction set. Therefore, application container images must be available for ARM64 before workloads can run on Graviton worker nodes. This can be achieved by rebuilding the images for ARM64 or, where appropriate, publishing multi-architecture images that support both x86_64 and ARM64 during migration.

After compatible images are deployed, EKS workloads can be scheduled onto Graviton node groups. AWS states that Graviton-based instances provide better price performance for many workloads and are generally priced lower than comparable x86 instances. Once the organization has established a predictable baseline of Graviton usage, a Compute Savings Plan can provide discounted pricing while retaining flexibility across eligible instance families, sizes, operating systems, and Regions. An EC2 Instance Savings Plan can also be appropriate when the instance family, Region, and operating system are sufficiently stable. See Amazon EKS cost optimization best practices and AWS Savings Plans documentation.

QUESTION NO: 149

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company uses AWS Control Tower for governance and uses AWS Transit Gateway for VPC connectivity across accounts.

In an AWS application account, the company's application team has deployed a web application that uses AWS Lambda and Amazon RDS. The company's database administrators have a separate DBA account and use the account to centrally manage all the databases across the organization. The database administrators use an Amazon EC2 instance that is deployed in the DBA account to access an RDS database that is deployed in the application account.

The application team has stored the database credentials as secrets in AWS Secrets Manager in the application account. The application team is manually sharing the secrets with the database administrators. The secrets are encrypted by the default AWS managed key for Secrets Manager in the application account. A solutions architect needs to implement a solution that gives the database administrators access to the database and eliminates the need to manually share the secrets.

Which solution will meet these requirements?

A. Use AWS Resource Access Manager (AWS RAM) to share the secrets from the application account with the DBA account. In the DBA account, create an IAM role that is named DBA-Admin. Grant the role the required permissions to access the shared secrets. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

B. In the application account, create an IAM role named DBA-Secret with permissions to access the secrets and a trust policy that allows the DBA-Admin role in the DBA account to assume it. In the DBA account, create the DBA-Admin role with permission to assume DBA-Secret, and attach DBA-Admin to the EC2 instance.

C. In the DBA account, create an IAM role that is named DBA-Admin. Grant the role the required permissions to access the secrets and the default AWS managed key in the application account. In the application account, attach resource-based policies to the key to allow access from the DBA account. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

D. In the DBA account, create an IAM role that is named DBA-Admin. Grant the role the required permissions to access the secrets in the application account. Attach an SCP to the application account to allow access to the secrets from the DBA account. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

ANSWER: B

Explanation:

Create the `DBA-Secret` role in the application account, authorize it to retrieve the required secret, and configure its trust policy to trust the `DBA-Admin` role in the DBA account. The EC2 instance receives the `DBA-Admin` role through its instance profile. That role can call AWS STS `AssumeRole` for `DBA-Secret`, then use the resulting temporary credentials to call Secrets Manager in the application account. This removes manual credential sharing and provides centrally managed, short-lived access that can be restricted to only the specific database secret and audited through CloudTrail.

Crucially, the assumed role is an identity in the application account, where both the secret and its default AWS managed Secrets Manager KMS key reside. AWS managed KMS keys cannot be modified to grant cross-account principals access. Using an assumed role in the owning account therefore permits Secrets Manager to decrypt and return the secret under the normal same-account authorization model, without requiring a customer managed KMS key or direct cross-account KMS permissions.

A role trust policy and the caller role's `sts:AssumeRole` permission establish the required cross-account delegation. AWS documents this model in IAM cross-account role access and documents the KMS requirements for cross-account access to Secrets Manager secrets.

QUESTION NO: 150

A team of data scientists is using Amazon SageMaker instances and SageMaker APIs to train machine learning (ML) models. The SageMaker instances are deployed in a

VPC that does not have access to or from the internet. Datasets for ML model training are stored in an Amazon S3 bucket. Interface VPC endpoints provide access to Amazon S3 and the SageMaker APIs.

Occasionally, the data scientists require access to the Python Package Index (PyPI) repository to update Python packages that they use as part of their workflow. A solutions architect must provide access to the PyPI repository while ensuring that the SageMaker instances remain isolated from the internet.

Which solution will meet these requirements?

A. Create an AWS CodeCommit repository for each package that the data scientists need to access. Configure code synchronization between the PyPI repositoryand the CodeCommit repository. Create a VPC endpoint for CodeCommit.

B. Create a NAT gateway in the VPC. Configure VPC routes to allow access to the internet with a network ACL that allows access to only the PyPI repositoryendpoint.

C. Create a NAT instance in the VPC. Configure VPC routes to allow access to the internet. Configure SageMaker notebook instance firewall rules that allow access to only the PyPI repository endpoint.

D. Create an AWS CodeArtifact domain and repository. Add an external connection for public:pypi to the CodeArtifact repository. Configure the Python client touse the CodeArtifact repository. Create a VPC endpoint for CodeArtifact.

ANSWER: D

Explanation:

Creating an AWS CodeArtifact domain and repository, adding an external connection for `public:pypi`, configuring the Python client to use that repository, and creating a VPC endpoint for CodeArtifact meets both the package-access and network-isolation requirements. CodeArtifact is a managed artifact repository that can proxy packages from supported public repositories, including PyPI, through an external connection. When a requested package is not already retained in the CodeArtifact repository, CodeArtifact retrieves it from PyPI and makes it available to authorized clients. This gives the data scientists a controlled internal package source rather than requiring their SageMaker resources to contact PyPI directly.

An interface VPC endpoint for CodeArtifact enables private connectivity from resources in the VPC to the CodeArtifact service by using AWS PrivateLink. The SageMaker instances therefore use private VPC networking and do not need an internet gateway, NAT device, or public IP address. Access can also be governed through IAM permissions and repository resource policies, allowing the organization to control which repositories and packages users can consume. The Python tooling is configured with the CodeArtifact repository endpoint and an authorization token so that normal package installation workflows retrieve packages through CodeArtifact. See the AWS CodeArtifact external connection documentation and the AWS CodeArtifact VPC endpoint documentation.

QUESTION NO: 151

A company runs a serverless application in a single AWS Region. The application accesses external URLs and extracts metadata from those sites. The company uses an Amazon Simple Notification Service (Amazon SNS) topic to publish URLs to an Amazon Simple Queue Service (Amazon SQS) queue An AWS Lambda function uses the queue as an event source and processes the URLs from the queue Results are saved to an Amazon S3 bucket

The company wants to process each URL other Regions to compare possible differences in site localization URLs must be published from the existing Region. Results must be written to the existing S3 bucket in the current Region.

Which combination of changes will produce multi-Region deployment that meets these requirements? (Select TWO.)

A. Deploy the SQS queue with the Lambda function to other Regions.

B. Subscribe the SNS topic in each Region to the SQS queue.

C. Subscribe the SQS queue in each Region to the SNS topics in each Region.

D. Configure the SQS queue to publish URLs to SNS topics in each Region.

E. Deploy the SNS topic and the Lambda function to other Regions.

F. Subscribe the SQS queue in each Region to the SNS topic in the existing Region.

ANSWER: A F

Explanation:

Deploying an Amazon SQS queue and its consuming Lambda function in each additional Region creates regional processing workers. This matters because Lambda event source mappings for Amazon SQS are regional: the Lambda function polls a queue in its own Region. Processing the same URL independently from each Region makes the outbound request originate from each chosen AWS Region, enabling the company to capture localization-dependent content or metadata.

Subscribing each regional SQS queue to the existing Region's SNS topic preserves the centralized publishing requirement while distributing every published URL to all regional workers. Amazon SNS supports cross-Region delivery to SQS, so the producer continues publishing once to the existing topic and SNS fans out a copy of the message to every subscribed queue. Each regional Lambda function can then process its queue's message and write results to the existing S3 bucket. The Lambda execution roles must grant the required S3 write permissions, and the bucket policy must allow those writes where applicable. This architecture separates centralized message publication from geographically distributed URL retrieval without requiring publishers or SNS topics in every processing Region.

See Amazon SNS cross-Region delivery and Using Lambda with Amazon SQS.

QUESTION NO: 152 - (SIMULATION)

A research company conducts mathematical simulations in the AWS Cloud. The simulations run on several hundred Amazon EC2 Linux instances. If a simulation encounters an issue, an engineer establishes an SSH connection to the affected EC2 instance.

Company policy requires that each EC2 instance session is established with a unique SSH key and that all SSH connections are logged in AWS CloudTrail. CloudTrail is enabled for the company's account and the AWS Regions that the company uses.

Which solution will meet these requirements?

ANSWER: See the explanation for the answer

Explanation:

Use Amazon EC2 Instance Connect (EIC) with a newly generated ephemeral SSH key pair for every connection.

Install/configure the EC2 Instance Connect package on the Linux EC2 instances (or use an AMI that includes it).
Grant engineers IAM permission to call `ec2-instance-connect:SendSSHPublicKey` only for authorized instances, OS users, and Availability Zones.
For each SSH session, the engineer generates a new SSH key pair and sends its public key to the target instance by using EC2 Instance Connect, for example:

```
aws ec2-instance-connect send-ssh-public-key \
  --instance-id i-0123456789abcdef0 \
  --availability-zone us-east-1a \
  --instance-os-user ec2-user \
  --ssh-public-key file://session-key.pub
```

The engineer then connects using the corresponding private key:

```
ssh -i session-key ec2-user@<instance-address>
```

EC2 Instance Connect injects the submitted public key temporarily (normally for 60 seconds), so a distinct key can be used for every session rather than maintaining shared or long-lived keys. The `SendSSHPublicKey` API call is recorded by AWS CloudTrail, providing the required audit trail for each SSH access attempt.

Do not rely solely on standard SSH key pairs stored in `authorized_keys`: SSH network traffic and ordinary SSH authentication events are not CloudTrail API events. EC2 Instance Connect supplies the CloudTrail-logged API action associated with the temporary key injection.

QUESTION NO: 153

A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.

All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX), and must accommodate high levels of throughput.

Which storage solution will meet these requirements?

A. Provision an AWS Storage Gateway file gateway NFS file share that is attached to an Amazon S3 bucket. Mount the NFS file share on each EC2 instance In the cluster.

B. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses General Purpose performance mode. Mount the EFS file system on each EC2 instance in the cluster.

C. Provision a new Amazon Elastic Block Store (Amazon EBS) volume that uses the lo2 volume type. Attach the EBS volume to all of the EC2 instances in the cluster.

D. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode. Mount the EFS file system on each EC2 instance in the cluster.

ANSWER: D

Explanation:

Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode. Mount the EFS file system on each EC2 instance in the cluster. Amazon EFS is a fully managed shared file system for Linux workloads. It can be mounted concurrently by many EC2 instances through the NFS protocol, providing all cluster nodes with shared read and write access to the same directory structure. EFS supports POSIX-compliant file-system semantics and permissions, which are required for applications that expect standard Linux file behavior.

An EFS file system configured for Regional availability stores data redundantly across multiple Availability Zones, providing the high availability and resilience required by a cluster distributed across Availability Zones. The Max I/O performance mode is designed for highly parallel workloads with a large number of clients and high aggregate throughput requirements, including big data analytics. This combination lets the cluster scale its shared storage access without requiring individual nodes to manage storage replication, failover, or file-system availability. AWS describes EFS shared access and regional resilience in the Amazon EFS User Guide, and documents performance-mode characteristics in the EFS performance documentation.

QUESTION NO: 154

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 CloudFormation 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 CloudFormation 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:

Developing infrastructure services using AWS CloudFormation templates and uploading each template as an AWS Service Catalog product to portfolios created in a central AWS account provides the central governance model required. Service Catalog administrators define the approved infrastructure products, versions, constraints, and access centrally. Portfolios can then be shared with an AWS Organization, organizational units, or individual member accounts, allowing the same approved catalog offerings to be distributed at scale without granting users authority to create arbitrary infrastructure outside the catalog.

Allowing user IAM roles to have ServiceCatalogEndUserAccess permissions only ensures that users can browse, launch, and manage their provisioned products through Service Catalog rather than directly calling broadly privileged infrastructure APIs. In recipient accounts, automation can import the shared portfolios, associate local principals, copy the required TagOptions, and configure launch constraints. A launch constraint causes Service Catalog to provision resources by using a predefined IAM role, separating end-user access from the permissions needed to create the underlying AWS resources.

Using the AWS Service Catalog TagOption Library supplies a controlled tag-key and tag-value taxonomy. Associating TagOptions with products or portfolios applies those governed tags during product provisioning, supporting the company's mandatory tagging requirement. See Sharing AWS Service Catalog portfolios and AWS Service Catalog TagOptions.

QUESTION NO: 155

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:

Provisioning a Direct Connect gateway, creating a private virtual interface on each Direct Connect connection, and associating the Direct Connect gateway with the VPC provides both required resiliency and scalable multi-Region connectivity. A Direct Connect gateway is a global AWS resource that lets on-premises networks reach VPCs in multiple AWS Regions through private virtual interfaces connected to the gateway. The two private virtual interfaces, one per Direct Connect connection, provide redundant paths between the on-premises network and AWS. BGP is used across the virtual interfaces, allowing routing to prefer an available path and maintain connectivity if one Direct Connect connection becomes unavailable.

For future expansion, additional VPCs in other Regions can be associated with the same Direct Connect gateway, subject to Direct Connect gateway association and routing limits. This avoids deploying separate Direct Connect connections for each Region and maintains private connectivity from the company network to supported VPC CIDR ranges. The existing private virtual interface must be replaced because a private VIF connected directly to a virtual private gateway cannot simply be moved to a Direct Connect gateway; a private VIF is created with its Direct Connect gateway attachment configuration. See the AWS Direct Connect gateway documentation and AWS private virtual interface documentation.

QUESTION NO: 156

A company is running several applications in the AWS Cloud. The applications are specific to separate business units in the company. The company is running the components of the applications in several AWS accounts that are in an organization in AWS Organizations. Every cloud resource in the company ' s organization has a tag that is named BusinessUnit. Every tag already has the appropriate value of the business unit name. The company needs to allocate its cloud costs to different business units. The company also needs to visualize the cloud costs for each business unit. Which solution will meet these requirements?

A. In the organization's management account, activate the existing BusinessUnit tag as a cost allocation tag. Also in the management account, create an Amazon S3 bucket and an AWS Cost and Usage Report (AWS CUR). Configure the S3 bucket as the destination for the AWS CUR. From the management account, query the AWS CUR data by using Amazon Athena. Use Amazon QuickSight for visualization.

B. In each member account, create a cost allocation tag that is named BusinessUnit. In the organization ' s management account, create an Amazon S3 bucket and an AWS Cost and Usage Report (AWS CUR). Configure the S3 bucket as the destination for the AWS CUR. Create an Amazon CloudWatch dashboard for visualization.

C. In the organization ' s management account, create a cost allocation tag that is named BusinessUnit. In each member account, create an Amazon S3 bucket and an AWS Cost and Usage Report (AWS CUR). Configure each S3 bucket as the destination for its respective AWS CUR. In the management account, create an Amazon CloudWatch dashboard for visualization.

D. In each member account, create a cost allocation tag that is named BusinessUnit. Also in each member account, create an Amazon S3 bucket and an AWS Cost and Usage Report (AWS CUR). Configure each S3 bucket as the destination for its respective AWS CUR. From the management account, query the AWS CUR data by using Amazon Athena. Use Amazon QuickSight for visualization.

ANSWER: A

Explanation:

Activate the existing user-defined `BusinessUnit` tag as a cost allocation tag in the AWS Organizations management account, then use a consolidated AWS Cost and Usage Report (CUR) as the central cost dataset. The organization’s management account has access to consolidated billing data for member accounts, so one CUR can capture the detailed cost and usage records needed across the organization. Once AWS activates the tag, its values are available in applicable billing data, enabling allocation, filtering, and grouping of costs by business unit. Activation affects data generated after activation; it does not retroactively add tag values to historical cost records.

Delivering the CUR to an Amazon S3 bucket provides durable centralized storage for the billing files. Amazon Athena can query CUR data in S3 with SQL, including the activated cost-allocation-tag columns, without requiring data to be loaded into a separate database. Amazon QuickSight can use the Athena query results as a dataset to build interactive dashboards, charts, and per-business-unit spend views. This combination provides both organization-wide cost allocation and flexible visualization at the detailed resource and line-item level. See the AWS documentation for activating user-defined cost allocation tags and querying CUR data with Athena.

QUESTION NO: 157

A greeting card company recently advertised that customers could send cards to their favourite celebrities through the company's platform Since the advertisement was published, the platform has received constant traffic from 10.000 unique users each second.

The platform runs on m5.xlarge Amazon EC2 instances behind an Application Load Balancer (ALB) The instances run in an Auto Scaling group and use a custom AMI that is based on Amazon Linux. The platform uses a highly available Amazon Aurora MySQL DB cluster that uses primary and reader endpoints The platform also uses an Amazon ElastiCache for Redis cluster that uses its cluster endpoint

The platform generates a new process for each customer and holds open database connections to MySQL for the duration of each customer's session However, resource usage for the platform is low.

Many customers are reporting errors when they connect to the platform Logs show that connections to the Aurora database are failing Amazon CloudWatch metrics show that the CPU load is tow across the platform and that connections to the platform are successful through the ALB.

Which solution will remediate the errors MOST cost-effectively?

A. Set up an Amazon CloudFront distribution Set the ALB as the origin Move all customer traffic to the CloudFront distribution endpoint

B. Use Amazon RDS Proxy Reconfigure the database connections to use the proxy

C. Increase the number of reader nodes in the Aurora MySQL cluster

D. Increase the number of nodes in the ElastiCache for Redis cluster

ANSWER: B

Explanation:

Use Amazon RDS Proxy Reconfigure the database connections to use the proxy is the most cost-effective solution because the failure pattern indicates Aurora MySQL connection exhaustion. The application creates a dedicated process for every customer and retains a database connection throughout each customer session. With sustained traffic of 10,000 unique users per second, the number of concurrently open connections can rapidly exceed the Aurora cluster's practical connection capacity, even though database and application CPU utilization remain low. Successful connections through the Application Load Balancer further isolate the issue to the database connectivity layer rather than frontend request distribution.

Amazon RDS Proxy sits between the application and an Aurora MySQL cluster and pools established database connections. It multiplexes and reuses those database connections across application requests, reducing the connection overhead and limiting the number of physical connections Aurora must maintain. The EC2 application should be reconfigured to use the proxy endpoint instead of connecting directly to the Aurora primary or reader endpoints. This addresses the immediate connection pressure without requiring larger database instances or architectural changes, and it can also improve resilience during database failovers. See the Amazon RDS Proxy User Guide and Amazon Aurora MySQL management documentation.