



"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

[www.atmicnetworks .com](http://www.atmicnetworks.com)

Warning: Keep connected with our support team
for latest updates

Question: 1

A company uses a service to collect metadata from applications that the company hosts on premises. Consumer devices such as TVs and internet radios access the applications. Many older devices do not support certain HTTP headers and exhibit errors when these headers are present in responses. The company has configured an on-premises load balancer to remove the unsupported headers from responses sent to older devices, which the company identified by the User-Agent headers.

The company wants to migrate the service to AWS, adopt serverless technologies, and retain the ability to support the older devices. The company has already migrated the applications into a set of AWS Lambda functions.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront distribution for the metadata service. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB. Configure the ALB to invoke the correct Lambda function for each type of request. Create a CloudFront function to remove the problematic headers based on the value of the User-Agent header.
- B. Create an Amazon API Gateway REST API for the metadata service. Configure API Gateway to invoke the correct Lambda function for each type of request. Modify the default gateway responses to remove the problematic headers based on the value of the User-Agent header.
- C. Create an Amazon API Gateway HTTP API for the metadata service. Configure API Gateway to invoke the correct Lambda function for each type of request. Create a response mapping template to remove the problematic headers based on the value of the User-Agent. Associate the response data mapping with the HTTP API.
- D. Create an Amazon CloudFront distribution for the metadata service. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB. Configure the ALB to invoke the correct Lambda function for each type of request. Create a Lambda@Edge function that will remove the problematic headers in response to viewer requests based on the value of the User-Agent header.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html>

Question: 2

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

Which solution will meet these requirements MOST cost-effectively?

- A. Upload the container images to AWS Lambda as functions. Configure a concurrency limit for the associated

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

Answer: B

Explanation:

minimizes operational + microservices that run on containers = AWS Elastic Beanstalk

Question: 3

A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application's data

a. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record.

The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an active-active strategy.

What should a solutions architect recommend to meet these requirements?

A. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region. Configure the CloudWatch alarm to invoke the Lambda function.

B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealthy. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.

C. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Region. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs. Remove the read replica. Replace the read replica with a standalone

RDS DB instance. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.

D. Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted targets. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region. Configure the CloudWatch alarm to invoke the Lambda function.

Answer: B

Explanation:

an AWS Lambda function in the backup region to promote the read replica and modify the Auto Scaling group values, and then configuring Route 53 with a health check that monitors the web application and sends an Amazon SNS notification to the Lambda function when the health check status is unhealthy. Finally, the application's Route 53 record should be updated with a failover policy that routes traffic to the ALB in the backup region when a health check failure occurs. This approach provides automatic failover to the backup region when a health check failure occurs, reducing the RTO to less than 15 minutes. Additionally, this approach is cost-effective as it does not require an active-active strategy.

Question: 4

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

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

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

- A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.
- B. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are configured in unlimited mode.
- C. Modify the DB instance to create a read replica in the same Availability Zone. Promote the read replica to be the primary DB instance in failure scenarios.
- D. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.
- E. Create a replication group for the ElastiCache for Redis cluster. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.
- F. Create a replication group for the ElastiCache for Redis cluster. Enable Multi-AZ on the cluster.

Answer: A,D,F

Explanation:

Option A is correct because using an Elastic Load Balancer and an Auto Scaling group with a minimum capacity

of two instances can improve the availability and scalability of the EC2 instances that host the application. The load balancer can distribute traffic across multiple instances and the Auto Scaling group can replace any unhealthy instances automatically¹

Option D is correct because modifying the DB instance to create a Multi-AZ deployment that extends across two Availability Zones can improve the availability and durability of the RDS for MariaDB database. Multi-AZ deployments provide enhanced data protection and minimize downtime by automatically failing over to a standby replica in another Availability Zone in case of a planned or unplanned outage⁴

Option F is correct because creating a replication group for the ElastiCache for Redis cluster and enabling Multi-AZ on the cluster can improve the availability and fault tolerance of the in-memory data store. A replication group consists of a primary node and up to five read-only replica nodes that are synchronized with the primary node using asynchronous replication. Multi-AZ allows automatic failover to one of the replicas if the primary node fails or becomes unreachable⁶

Reference: 1: <https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/how-elastic-load-balancing-works.html> 2: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/burstable-performance-instances-unlimited-mode.html> 3:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html 4:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html> 5:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoScaling.html> 6:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.Redis.Groups.html>

Question: 5

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 a custom error response by configuring a CloudFront custom error page. Modify DNS records to point to a publicly accessible web page.

Answer: C,E

Explanation:

"Save your custom error pages in a location that is accessible to CloudFront. We recommend that you store them in an Amazon S3 bucket, and that you don't store them in the same place as the rest of your website or application's content. If you store the custom error pages on the same origin as your website or application, and the origin starts to return 5xx errors, CloudFront can't get the custom error pages because the origin server is

unavailable." <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.html>

Question: 6

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: A,E

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html>

Question: 7

A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC. The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege. Which solution meets these requirements?

- A. Create an AWS PrivateLink interface VPC endpoint. Connect this endpoint to the endpoint service that the third-party SaaS application provides. Create a security group to limit the access to the endpoint. Associate the security group with the endpoint.
- B. Create an AWS Site-to-Site VPN connection between the third-party SaaS application and the company VPC. Configure network ACLs to limit access across the VPN tunnels.
- C. Create a VPC peering connection between the third-party SaaS application and the company VPC. Update route tables by adding the needed routes for the peering connection.
- D. Create an AWS PrivateLink endpoint service. Ask the third-party SaaS provider to create an interface VPC endpoint for this endpoint service. Grant permissions for the endpoint service to the specific account of the third-party SaaS provider.

Answer: A

Explanation:

Reference architecture - <https://docs.aws.amazon.com/vpc/latest/privatelink/privatelink-access-saas.html>
Note from documentation that Interface Endpoint is at client side

Question: 8

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:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html>

Question: 9

A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.

Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instance. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group. Invoke an AWS Lambda function on the `autoscaling:EC2_INSTANCE_TERMINATING` transition to send `ABANDON` to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.

B. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group. Invoke an AWS Lambda function on the `autoscaling:EC2_INSTANCE_TERMINATING` transition to call the AWS Systems Manager API

`SendCommand` operation to run the document to copy the log files and send `CONTINUE` to the Auto Scaling group to terminate the instance.

C. Change the log delivery rate to every 5 minutes. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user data. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance termination. Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.

D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic. From the SNS notification, call the AWS Systems Manager API `SendCommand` operation to run the document to copy the log files and send `ABANDON` to the Auto Scaling group to terminate the instance.

Answer: B

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/adding-lifecycle-hooks.html>

- Refer to Default Result section - If the instance is terminating, both `abandon` and `continue` allow the instance to terminate. However, `abandon` stops any remaining actions, such as other lifecycle hooks, and `continue`

allows any other lifecycle hooks to complete.

<https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-instance/>

<https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function> <https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function/blob/master/cloudformation/template.yaml>

Question: 10

A company is using multiple AWS accounts. The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A. The company's applications and databases are running in Account B. A solutions architect will deploy a two-net application in a new VPC. To simplify the configuration, the db.example.com CNAME record set for the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53. During deployment, the application failed to start. Troubleshooting revealed that db.example.com is not resolvable on the Amazon EC2 instance. The solutions architect confirmed that the record set was created correctly in Route 53.

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

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

Answer: C,E

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/private-hosted-zone-different-account/>

Question: 11

A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto Scaling group. The web application stores all blog content on an Amazon EFS volume.

The company recently added a feature for bloggers to add video to their posts, attracting 10 times the previous user traffic. At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.

Which is the MOST cost-efficient and scalable deployment that will resolve the issues for users?

- A. Reconfigure Amazon EFS to enable maximum I/O.
- B. Update the Nog site to use instance store volumes for storage. Copy the site contents to the volumes at launch and to Amazon S3 at shutdown.
- C. Configure an Amazon CloudFront distribution. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
- D. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/>

Using an Amazon S3 bucket

Using a MediaStore container or a MediaPackage channel

Using an Application Load Balancer

Using a Lambda function URL

Using Amazon EC2 (or another custom origin)

Using CloudFront origin groups

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

Question: 12

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:

A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions. The following describe scenarios where you can use a Direct Connect gateway. <https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html>

Question: 13

A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.

The website contains static content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2 instances running in an Auto Scaling group to process an Amazon SQS queue. The company wants to re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.

Which solution meets these requirements?

- A. Use Amazon ECS containers for the web application and Spot Instances for the Auto Scaling group that processes the SQS queue. Replace the custom software with Amazon Rekognition to categorize the videos.
- B. Store the uploaded videos in Amazon EFS and mount the file system to the EC2 instances for the web application. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
- C. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notifications to publish events to the SQS queue. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
- D. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the web application and launch a worker environment to process the SQS queue. Replace the custom software with Amazon Rekognition to categorize the videos.

Answer: C

Explanation:

Option C is correct because hosting the web application in Amazon S3, storing the uploaded videos in Amazon S3, and using S3 event notifications to publish events to the SQS queue reduces the operational overhead of managing EC2 instances and EBS volumes. Amazon S3 can serve static content such as HTML, CSS, JavaScript, and media files directly from S3 buckets. Amazon S3 can also trigger AWS Lambda functions through S3 event notifications when new objects are created or existing objects are updated or deleted. AWS Lambda can process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize

the videos. This solution eliminates the need for custom recognition software and third-party dependencies³⁴⁵
Reference: 1: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html> 2:
<https://aws.amazon.com/efs/pricing/> 3:<https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html> 4:
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/NotificationHowTo.html> 5:
<https://docs.aws.amazon.com/rekognition/latest/dg/what-is.html> 6:<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html>

Question: 14

A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.

How can this be accomplished?

- A. Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda function. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.
- B. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify code. Rollback if Amazon CloudWatch alarms are triggered.
- C. Refactor the AWS CLI scripts into a single script that deploys the new Lambda version. When deployment is completed, the script tests execute. If errors are detected, revert to the previous Lambda version.
- D. Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda version. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

Answer: B

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and-phased-deployments-with-aws-codedeploy/>

Question: 15

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.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/aws/new-attributes-based-access-control-with-aws-single-sign-on/>

Question: 16

A software company has deployed an application that consumes a REST API by using Amazon API Gateway, AWS Lambda functions, and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.

A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.

What should the solutions architect recommend to improve the customer experience?

- A. Implement retry logic with exponential backoff and irregular variation in the client application. Ensure that the errors are caught and handled with descriptive error messages.
- B. Implement API throttling through a usage plan at the API Gateway level. Ensure that the client application handles code 429 replies without error.
- C. Turn on API caching to enhance responsiveness for the production stage. Run 10-minute load tests. Verify that the cache capacity is appropriate for the workload.
- D. Implement reserved concurrency at the Lambda function level to provide the resources that are needed during sudden increases in traffic.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/aws-batch-requests-error/>
<https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-429-limit/>

Question: 17

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

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

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

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

Answer: A

Explanation:

<https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and-amazon-s3/>

Question: 18

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 EC2 instances 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 companies to add to their allow lists.
- 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 other companies to add to their allow lists.

Answer: C

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html> 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. As it uses the NLB as the resource in the A-record, traffic will be routed through the NLB, and it will automatically route the traffic to the healthy instances based on the health checks and also it provides the fixed address assignments as the other companies can add the NLB's Elastic IP addresses to their allow lists.

Question: 19

A company uses an on-premises data analytics platform. The system is highly available in a fully redundant configuration across 12 servers in the company's data center.

The system runs scheduled jobs, both hourly and daily, in addition to one-time requests from users. Scheduled jobs can take between 20 minutes and 2 hours to finish running and have tight SLAs. The scheduled jobs account for 65% of the system usage. User jobs typically finish running in less than 5 minutes and have no SL

- A. The user jobs account for 35% of system usage. During system failures, scheduled jobs must continue to

meet SLAs. However, user jobs can be delayed.

A solutions architect needs to move the system to Amazon EC2 instances and adopt a consumption-based model to reduce costs with no long-term commitments. The solution must maintain high availability and must not affect the SLAs.

Which solution will meet these requirements MOST cost-effectively?

- A. Split the 12 instances across two Availability Zones in the chosen AWS Region. Run two instances in each Availability Zone as On-Demand Instances with Capacity Reservations. Run four instances in each Availability Zone as Spot Instances.
- B. Split the 12 instances across three Availability Zones in the chosen AWS Region. In one of the Availability Zones, run all four instances as On-Demand Instances with Capacity Reservations. Run the remaining instances as Spot Instances.
- C. Split the 12 instances across three Availability Zones in the chosen AWS Region. Run two instances in each Availability Zone as On-Demand Instances with a Savings Plan. Run two instances in each Availability Zone as Spot Instances.
- D. Split the 12 instances across three Availability Zones in the chosen AWS Region. Run three instances in each Availability Zone as On-Demand Instances with Capacity Reservations. Run one instance in each Availability Zone as a Spot Instance.

Answer: D

Explanation:

By splitting the 12 instances across three Availability Zones, the system can maintain high availability and availability of resources in case of a failure. Option D also uses a combination of On-Demand Instances with Capacity Reservations and Spot Instances, which allows for scheduled jobs to be run on the On-Demand instances with guaranteed capacity, while also taking advantage of the cost savings from Spot Instances for the user jobs which have lower SLA requirements.

Question: 20

A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an encrypted file in Amazon S3. For the next version of the application, the security engineer wants to implement the following application design changes to improve security: The database must use strong, randomly generated passwords stored in a secure AWS managed service.

The application resources must be deployed through AWS CloudFormation.

The application must rotate credentials for the database every 90 days.

A solutions architect will generate a CloudFormation template to deploy the application. Which resources specified in the CloudFormation template will meet the security engineer's requirements with the LEAST amount of operational overhead?

- A. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Specify a Secrets Manager RotationSchedule resource to

rotate the database password every 90 days.

B. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Store. Create an AWS Lambda function resource to rotate the database password. Specify a Parameter Store RotationSchedule resource to rotate the database password every 90 days.

C. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Create an Amazon EventBridge scheduled rule resource to trigger the Lambda function password rotation every 90 days.

D. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Store. Specify an AWS AppSync DataSource resource to automatically rotate the database password every 90 days.

Answer: B

Explanation: <https://aws.amazon.com/blogs/security/how-to-securely-provide-database-credentials-to-lambda-functions-by-using-aws-secrets-manager/>
<https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>

https://docs.aws.amazon.com/secretsmanager/latest/userguide/integrating_cloudformation.html

Question: 21

A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand.

Which solutions meet these requirements? (Choose two.)

A. Create an Amazon API Gateway REST API. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.

B. Create an Amazon API Gateway HTTP API. Configure this API with direct integrations to Dynamo DB by using API Gateway's AWS integration type.

C. Create an Amazon API Gateway HTTP API. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.

D. Create an accelerator in AWS Global Accelerator. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.

E. Create a Network Load Balancer. Configure listener rules to forward requests to the appropriate AWS Lambda functions

Answer: A,C

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-overview-developer-experience.html>

Question: 22

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

A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests. Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

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

Answer: C,E,F

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.html>

Question: 23

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

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

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

What should a solutions architect do to meet these requirements?

- A. In the management account of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.
- B. In the member accounts of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Schedule a monthly AWS Lambda function to retrieve the reports and calculate the total cost for the costCenter tagged resources.
- C. In the member accounts of the organization activate the costCenter user-defined tag. From the management account, schedule a monthly AWS Cost and Usage Report. Use the tag breakdown in the report to calculate the total cost for the costCenter tagged resources.
- D. Create a custom report in the organization view in AWS Trusted Advisor. Configure the report to generate a monthly billing summary for the costCenter tagged resources in the compliance team's AWS account.

Answer: A

Explanation:

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html>
<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/configurecostallocreport.html>

Question: 24

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

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

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

Answer: A,C

Explanation:

<https://aws.amazon.com/blogs/mt/self-service-vpcs-in-aws-control-tower-using-aws-service-catalog/>
<https://docs.aws.amazon.com/vpc/latest/tgw/tgw-transit-gateways.html>
<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-transitgatewayattachment.html>

Question: 25

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

- A. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the PowerUserAccess managed policy to the role. Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPriateMarketplaceAdminFullAccess managed policy.
- B. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the AdministratorAccess managed policy to the role. Define a permissions boundary with the AWSPriateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.
- C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization. Add the AWSPriateMarketplaceAdminFullAccess managed policy to the role. Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.
- D. Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developers. Add the AWSPriateMarketplaceAdminFullAccess managed policy to the role. Create an SCP in Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Apply the SCP to all the shared services accounts in the organization.

Answer: C

Explanation:

SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. <https://aws.amazon.com/blogs/awsmarketplace/controlling-access-to-a-well-architected-private-marketplace-using-iam-and-aws-organizations/>

This approach allows the procurement managers to assume the procurement-manager-role in shared services accounts, which have the AWSPriateMarketplaceAdminFullAccess managed policy attached to it and can then manage the Private Marketplace. The organization root-level SCP denies the permission to administer Private Marketplace to everyone except the role named procurementmanager-role and another SCP denies the permission to create an IAM role named procurementmanager-role to everyone in the organization, ensuring that only the procurement team can assume the role and manage the Private Marketplace. This approach provides a centralized way to manage and restrict access to Private Marketplace while maintaining a high level

of security.

Question: 26

A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2, Amazon S3 and Amazon DynamoDB. The developers account resides in a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

```
"Version": "2012-10-17",
"statement": [
  {
    "Sid": "AilcwECS"*
    "Effect": "Allow",
    "Action": "ec2:*", "Resource": "*"
  },
  {
    "Sid": "AllOMDynamCDH^
    "Effect": "Allow",
    "Action": "dynamodb:*",
    "Resource": "*"
  },
  {
    "Sid": "Allcws",
    "Effect": "Allow",
    "Resource": "*"
  }
]
```

When this policy is deployed, IAM users in the developers account are still able to use AWS services that are not listed in the policy. What should the solutions architect do to eliminate the developers' ability to use services outside the scope of this policy?

- A. Create an explicit deny statement for each AWS service that should be constrained
- B. Remove the Full AWS Access SCP from the developer account's OU
- C. Modify the Full AWS Access SCP to explicitly deny all services
- D. Add an explicit deny statement using a wildcard to the end of the SCP

Answer: B

Explanation:

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

Question: 27

A company is hosting a monolithic REST-based API for a mobile app on five Amazon EC2 instances in public subnets of a VPC. Mobile clients connect to the API by using a domain name that is hosted on Amazon Route 53. The company has created a Route 53 multivalued answer routing policy with the IP addresses of all the EC2 instances. Recently, the app has been overwhelmed by large and sudden increases in traffic. The app has not been able to keep up with the traffic.

A solutions architect needs to implement a solution so that the app can handle the new and varying load.

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

- A. Separate the API into individual AWS Lambda functions. Configure an Amazon API Gateway REST API with Lambda integration for the backend. Update the Route 53 record to point to the API Gateway API.
- B. Containerize the API logic. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Run the containers in the cluster by using Amazon EC2. Create a Kubernetes ingress. Update the Route 53 record to point to the Kubernetes ingress.
- C. Create an Auto Scaling group. Place all the EC2 instances in the Auto Scaling group. Configure the Auto Scaling group to perform scaling actions that are based on CPU utilization. Create an AWS Lambda function that reacts to Auto Scaling group changes and updates the Route 53 record.
- D. Create an Application Load Balancer (ALB) in front of the API. Move the EC2 instances to private subnets in the VPC. Add the EC2 instances as targets for the ALB. Update the Route 53 record to point to the ALB.

Answer: D

Explanation:

By breaking down the monolithic API into individual Lambda functions and using API Gateway to handle the incoming requests, the solution can automatically scale to handle the new and varying load without the need for manual scaling actions. Additionally, this option will automatically handle the traffic without the need of having EC2 instances running all the time and only pay for the number of requests and the duration of the execution of the Lambda function.

By updating the Route 53 record to point to the API Gateway, the solution can handle the traffic and also it will direct the traffic to the correct endpoint.

Question: 28

A company has created an OU in AWS Organizations for each of its engineering teams. Each OU owns multiple AWS accounts. The organization has hundreds of AWS accounts. A solutions architect must design a solution so that each OU can view a breakdown of usage costs across its AWS accounts. Which solution meets these requirements?

- A. Create an AWS Cost and Usage Report (CUR) for each OU by using AWS Resource Access Manager Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
- B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account- Allow each team to visualize the CUR through an Amazon QuickSight dashboard
- C. Create an AWS Cost and Usage Report (CUR) in each AWS Organizations member account Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
- D. Create an AWS Cost and Usage Report (CUR) by using AWS Systems Manager Allow each team to visualize the CUR through Systems Manager OpsCenter dashboards

Answer: B

Explanation:

<https://docs.aws.amazon.com/cur/latest/userguide/billing-cur-limits.html>

Question: 29

A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.

The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS.

Which data migration strategy should the company use?

- A. Use the file gateway option in AWS Storage Gateway to replace the existing Windows file server, and point the existing file share to the new file gateway.
- B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.
- C. Use AWS Data Pipeline to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).
- D. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS),

Answer: B

Explanation: <https://aws.amazon.com/storagegateway/file/>

[https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-](https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html)

[datasync.htmlhttps://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os-windows-server](https://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os-windows-server)

Question: 30

A company's solutions architect is reviewing a web application that runs on AWS. The application references static assets in an Amazon S3 bucket in the us-east-1 Region. The company needs resiliency across multiple AWS

Regions. The company already has created an S3 bucket in a second Region.

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

- A. Configure the application to write each object to both S3 buckets. Set up an Amazon Route 53 public hosted zone with a record set by using a weighted routing policy for each S3 bucket. Configure the application to reference the objects by using the Route 53 DNS name.
- B. Create an AWS Lambda function to copy objects from the S3 bucket in us-east-1 to the S3 bucket in the second Region. Invoke the Lambda function each time an object is written to the S3 bucket in us-east-1. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
- C. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
- D. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region. If failover is required, update the application code to load S3 objects from the S3 bucket in the second Region.

Answer: C

Explanation:

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

Question: 31

A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database. A solutions architect must design a scalable and highly available solution to meet the demand of 200,000 daily users.

Which steps should the solutions architect take to design an appropriate solution?

- A. Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance. The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones. Use an Amazon Route 53 alias record to route traffic from the company's domain to the NLB.
- B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.
- C. Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Region. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica. Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.
- D. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot Instances spanning three Availability Zones. The stack should launch an Amazon RDS

MySQL DB instance with a Snapshot deletion policy Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB

Answer: C

Explanation:

Using AWS CloudFormation to launch a stack with an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones, a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy, and an Amazon Route 53 alias record to route traffic from the company's domain to the ALB will ensure that

Question: 32

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

What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

- A. Create a stack set in the Organizations member accounts. Use service-managed permissions. Set deployment options to deploy to an organization. Use CloudFormation StackSets drift detection.
- B. Create stacks in the Organizations member accounts. Use self-service permissions. Set deployment options to deploy to an organization. Enable the CloudFormation StackSets automatic deployment.
- C. Create a stack set in the Organizations management account Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets automatic deployment.
- D. Create stacks in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets drift detection.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-accounts-and-regions/>

Question: 33

A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises infrastructure that consists of physical machines and VMs that host numerous applications.

The company must capture details about the system configuration, system performance, running processes and network configurations of its on-premises servers. The company also must divide the on-premises

applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.

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

- A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs.
- B. Assess the existing applications by installing AWS Systems Manager Agent on the physical machines and VMs
- C. Group servers into applications for migration by using AWS Systems Manager Application Manager.
- D. Group servers into applications for migration by using AWS Migration Hub.
- E. Generate recommended instance types and associated costs by using AWS Migration Hub.
- F. Import data about server sizes into AWS Trusted Advisor. Follow the recommendations for cost optimization.

Answer: A,D,E

Explanation:

<https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html><https://docs.aws.amazon.com/migrationhub/latest/ug/ec2-recommendations.html>

Question: 34

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 11 GiB 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:

Create Amazon S3 gateway endpoint in the VPC and add a VPC endpoint policy. This VPC endpoint policy will have a statement that allows S3 access only via access points owned by the organization.

Question: 35

A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load for the rest of the week. The access pattern includes many more writes to the table than reads of the table.

A solutions architect needs to implement a solution to minimize the cost of the table.

Which solution will meet these requirements?

- A. Use AWS Application Auto Scaling to increase capacity during the peak period. Purchase reserved RCUs and WCUs to match the average load.
- B. Configure on-demand capacity mode for the table.
- C. Configure DynamoDB Accelerator (DAX) in front of the table. Reduce the provisioned read capacity to match the new peak load on the table.
- D. Configure DynamoDB Accelerator (DAX) in front of the table. Configure on-demand capacity mode for the table.

Answer: D

Explanation:

This solution meets the requirements by using Application Auto Scaling to automatically increase capacity during the peak period, which will handle the double the average load. And by purchasing reserved RCUs and WCUs to match the average load, it will minimize the cost of the table for the rest of the week when the load is close to the average.

Question: 36

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

What is the MOST cost-effective migration recommendation?

- A. Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. When

there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in an Amazon S3 bucket.

B. Create a queue using Amazon MQ. Configure the existing web server to publish to the new queue. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the files. Store the processed files in Amazon EFS. Shut down the EC2 instance after the task is complete.

C. Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in Amazon EFS.

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

Answer: D

Explanation:

<https://aws.amazon.com/blogs/compute/operating-lambda-performance-optimization-part-1/>

Question: 37

A company is using Amazon OpenSearch Service to analyze data.

a. The company loads data into an OpenSearch Service cluster with 10 data nodes from an Amazon S3 bucket that uses S3 Standard storage. The data resides in the cluster for 1 month for read-only analysis. After 1 month, the company deletes the index that contains the data from the cluster. For compliance purposes, the company must retain a copy of all input data.

The company is concerned about ongoing costs and asks a solutions architect to recommend a new solution.

Which solution will meet these requirements MOST cost-effectively?

A. Replace all the data nodes with UltraWarm nodes to handle the expected capacity. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

B. Reduce the number of data nodes in the cluster to 2. Add UltraWarm nodes to handle the expected capacity. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data. Transition the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy.

C. Reduce the number of data nodes in the cluster to 2. Add UltraWarm nodes to handle the expected capacity. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data. Add cold storage nodes to the cluster. Transition the indexes from UltraWarm to cold storage. Delete the input data from the S3 bucket after 1 month by using an S3 Lifecycle policy.

D. Reduce the number of data nodes in the cluster to 2. Add instance-backed data nodes to handle the expected capacity. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

Answer: B

Explanation:

By reducing the number of data nodes in the cluster to 2 and adding UltraWarm nodes to handle the expected capacity, the company can reduce the cost of running the cluster. Additionally, configuring the indexes to transition to UltraWarm when OpenSearch Service ingests the data will ensure that the data is stored in the most cost-effective manner. Finally, transitioning the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy will ensure that the data is retained for compliance purposes, while also reducing the ongoing costs.

Question: 38

A company has 10 accounts that are part of an organization in AWS Organizations. AWS Config is configured in each account. All accounts belong to either the Prod OU or the NonProd OU. The company has set up an Amazon EventBridge rule in each AWS account to notify an Amazon Simple Notification Service (Amazon SNS) topic when an Amazon EC2 security group inbound rule is created with 0.0.0.0/0 as the source. The company's security team is subscribed to the SNS topic. For all accounts in the NonProd OU, the security team needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source. Which solution will meet this requirement with the LEAST operational overhead?

- A. Modify the EventBridge rule to invoke an AWS Lambda function to remove the security group inbound rule and to publish to the SNS topic. Deploy the updated rule to the NonProd OU.
- B. Add the vpc-sg-open-only-to-authorized-ports AWS Config managed rule to the NonProd OU.
- C. Configure an SCP to allow the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is not 0.0.0.0/0. Apply the SCP to the NonProd OU.
- D. Configure an SCP to deny the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0. Apply the SCP to the NonProd OU.

Answer: D

Explanation:

This solution will meet the requirement with the least operational overhead because it directly denies the creation of the security group inbound rule with 0.0.0.0/0 as the source, which is the exact requirement. Additionally, it does not require any additional steps or resources such as invoking a Lambda function or adding a Config rule.

An SCP (Service Control Policy) is a policy that you can use to set fine-grained permissions for your AWS accounts within your organization. You can use SCPs to set permissions for the root user of an account and to delegate permissions to IAM users and roles in the accounts. You can use SCPs to set permissions that allow or deny access to specific services, actions, and resources.

To implement this solution, you would need to create an SCP that denies the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0. This SCP would then be applied to the NonProd OU. This would ensure that any security group inbound rule that includes 0.0.0.0/0 as the source will be denied, thus meeting the requirement.

Reference:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_condition-keys.html

Question: 39

A company hosts a Git repository in an on-premises data center. The company uses webhooks to invoke functionality that runs in the AWS Cloud. The company hosts the webhook logic on a set of Amazon EC2 instances in an Auto Scaling group that the company set as a target for an Application Load Balancer (ALB). The Git server calls the ALB for the configured webhooks. The company wants to move the solution to a serverless architecture.

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

- A. For each webhook, create and configure an AWS Lambda function URL. Update the Git servers to call the individual Lambda function URLs.
- B. Create an Amazon API Gateway HTTP API. Implement each webhook logic in a separate AWS Lambda function. Update the Git servers to call the API Gateway endpoint.
- C. Deploy the webhook logic to AWS App Runner. Create an ALB, and set App Runner as the target. Update the Git servers to call the ALB endpoint.
- D. Containerize the webhook logic. Create an Amazon Elastic Container Service (Amazon ECS) cluster, and run the webhook logic in AWS Fargate. Create an Amazon API Gateway REST API, and set Fargate as the target. Update the Git servers to call the API Gateway endpoint.

Answer: B

Explanation:

<https://aws.amazon.com/solutions/implementations/git-to-s3-using-webhooks/>
<https://medium.com/mindorks/building-webhook-is-easy-using-aws-lambda-and-api-gateway-56f5e5c3a596>

Question: 40

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:

it covers all the requirements mentioned in the question, it will allow collecting the detailed metrics, including process information and it provides a way to query and analyze the data using Amazon Athena.

Question: 41

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

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

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

- A. Deploy a NAT gateway. Associate an Elastic IP address with the NAT gateway. Configure the VPC to use the NAT gateway.
- B. Deploy an egress-only internet gateway. Associate an Elastic IP address with the egress-only internet gateway. Configure the elastic network interface on the Lambda function to use the egress-only internet gateway.
- C. Deploy an internet gateway. Associate an Elastic IP address with the internet gateway. Configure the Lambda function to use the internet gateway.
- D. Deploy an internet gateway. Associate an Elastic IP address with the internet gateway. Configure the default route in the public VPC route table to use the internet gateway.

Answer: A

Explanation:

This solution will give the Lambda function access to the internet by routing its outbound traffic through the NAT gateway, which has a public Elastic IP address. This will allow the external provider to whitelist the single public IP address associated with the NAT gateway, and enable the application to access the new service

Deploying a NAT gateway and associating an Elastic IP address with it, and then configuring the VPC to use the

NAT gateway, will give the application the ability to access the new service. This is because the NAT gateway will be the single public IP address that the external provider needs for the allow list. The NAT gateway will allow the application to access the service, while keeping the underlying Lambda functions private. When configuring NAT gateways, you should ensure that the route table associated with the NAT gateway has a route to the internet gateway with a target of the internet gateway. Additionally, you should ensure that the security group associated with the NAT gateway allows outbound traffic from the Lambda functions.

Reference:

AWS Certified Solutions Architect Professional Official Amazon Text Book[1], page 456

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Gateway.html

Question: 42

A solutions architect has developed a web application that uses an Amazon API Gateway Regional endpoint and an AWS Lambda function. The consumers of the web application are all close to the AWS Region where the application will be deployed. The Lambda function only queries an Amazon Aurora MySQL database. The solutions architect has configured the database to have three read replicas.

During testing, the application does not meet performance requirements. Under high load, the application opens a large number of database connections. The solutions architect must improve the application's performance.

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

- A. Use the cluster endpoint of the Aurora database.
- B. Use RDS Proxy to set up a connection pool to the reader endpoint of the Aurora database.
- C. Use the Lambda Provisioned Concurrency feature.
- D. Move the code for opening the database connection in the Lambda function outside of the event handler.
- E. Change the API Gateway endpoint to an edge-optimized endpoint.

Answer: B,D

Explanation:

Connect to RDS outside of Lambda handler method to improve

performance <https://awstut.com/en/2022/04/30/connect-to-rds-outside-of-lambda-handler-method-to-improve-performance-en/>

Using RDS Proxy, you can handle unpredictable surges in database traffic. Otherwise, these surges might cause issues due to oversubscribing connections or creating new connections at a fast rate. RDS Proxy establishes a database connection pool and reuses connections in this pool. This approach avoids the memory and CPU overhead of opening a new database connection each time. To protect the database against oversubscription, you can control the number of database connections that are created. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html>

Question: 43

A company is planning to host a web application on AWS and works to load balance the traffic across a group of

Amazon EC2 instances. One of the security requirements is to enable end-to-end encryption in transit between the client and the web server.

Which solution will meet this requirement?

- A. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Export the SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
- B. Associate the EC2 instances with a target group. Provision an SSL certificate using AWS Certificate Manager (ACM). Create an Amazon CloudFront distribution and configure it to use the SSL certificate. Set CloudFront to use the target group as the origin server.
- C. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Provision a third-party SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
- D. Place the EC2 instances behind a Network Load Balancer (NLB). Provision a third-party SSL certificate and install it on the NLB and on each EC2 instance. Configure the NLB to listen on port 443 and to forward traffic to port 443 on the instances.

Answer: A

Explanation:

Option A is correct because placing the EC2 instances behind an Application Load Balancer (ALB) and associating an SSL certificate from AWS Certificate Manager (ACM) with the ALB enables encryption in transit between the client and the ALB. Exporting the SSL certificate and installing it on each EC2 instance enables encryption in transit between the ALB and the web server. Configuring the ALB to listen on port 443 and to forward traffic to port 443 on the instances ensures that HTTPS is used for both connections. This solution achieves end-to-end encryption in transit for the web application¹²

Reference: 1: <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html> 2: <https://docs.aws.amazon.com/acm/latest/userguide/acm-overview.html> 3: <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html> : <https://aws.amazon.com/certificate-manager/faqs/> : <https://docs.aws.amazon.com/elasticloadbalancing/latest/network/introduction.html>

Question: 44

A company wants to migrate its data analytics environment from on premises to AWS. The environment consists of two simple Node.js applications. One of the applications collects sensor data and loads it into a MySQL database. The other application aggregates the data into reports. When the aggregation jobs run, some of the load jobs fail to run correctly.

The company must resolve the data loading issue. The company also needs the migration to occur without interruptions or changes for the company's customers.

What should a solutions architect do to meet these requirements?

- A. Set up an Amazon Aurora MySQL database as a replication target for the on-premises database. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind a Network Load Balancer (NLB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the NLB.
- B. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Move the aggregation jobs to run against the Aurora MySQL database. Set up collection endpoints behind an Application Load Balancer (ALB) as Amazon EC2 instances in an Auto Scaling group. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on-premises to AWS.
- C. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB) and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on-premises to AWS.
- D. Set up an Amazon Aurora MySQL database. Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as an Amazon Kinesis data stream. Use Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the Kinesis data stream.

Answer: C

Explanation:

Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on-premises to AWS.

Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%.

Question: 45

A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled.

Which solution will meet these requirements?

- A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed key. Use the AWS CLI to re-upload all objects in the S3 bucket. Set an S3 bucket policy to deny unencrypted PutObject requests.
- B. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject requests. Use the AWS CLI to re-upload all objects in the S3 bucket.
- C. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to automatically encrypt objects on GetObject and PutObject requests.
- D. In the S3 bucket properties, change the default encryption to AES-256 with a customer managed key. Attach a policy to deny unencrypted PutObject requests to any entities that access the S3 bucket. Use the AWS CLI to re-upload all objects in the S3 bucket.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerSideEncryptionCustomerKeys.html>
Clearly says we need following header for SSE-C x-amz-server-side-encryption-customer-algorithm
Use this header to specify the encryption algorithm. The header value must be AES256.

Question: 46

A company is running a web application in the AWS Cloud. The application consists of dynamic content that is created on a set of Amazon EC2 instances. The

EC2 instances run in an Auto Scaling group that is configured as a target group for an Application Load Balancer (ALB).

The company is using an Amazon CloudFront distribution to distribute the application globally. The CloudFront distribution uses the ALB as an origin. The company uses Amazon Route 53 for DNS and has created an A record of www.example.com for the CloudFront distribution.

A solutions architect must configure the application so that it is highly available and fault tolerant. Which solution meets these requirements?

- A. Provision a full, secondary application deployment in a different AWS Region. Update the Route 53 A record to be a failover record. Add both of the CloudFront distributions as values. Create Route 53 health checks.
- B. Provision an ALB, an Auto Scaling group, and EC2 instances in a different AWS Region. Update the CloudFront distribution, and create a second origin for the new ALB. Create an origin group for the two origins. Configure one origin as primary and one origin as secondary.
- C. Provision an Auto Scaling group and EC2 instances in a different AWS Region. Create a second target for the new Auto Scaling group in the ALB. Set up the failover routing algorithm on the ALB.
- D. Provision a full, secondary application deployment in a different AWS Region. Create a second CloudFront distribution, and add the new application setup as an origin. Create an AWS Global Accelerator accelerator. Add both of the CloudFront distributions as endpoints.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html>

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

You can set up CloudFront with origin failover for scenarios that require high availability. To get started, you create an origin group with two origins: a primary and a secondary. If the primary origin is unavailable, or returns specific HTTP response status codes that indicate a failure, CloudFront automatically switches to the secondary origin.

Question: 47

A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts. AWS Site-to-Site VPN connections are configured between all of the company's global offices and the transit account. The company has AWS Config enabled on all of its accounts.

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices. Developers will reference this list to gain access to applications securely.

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

- A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges. Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is updated. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with the updated IP address ranges.
- B. Create a new AWS Config managed rule that contains all of the internal IP address ranges. Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address ranges. Configure the rule to automatically remediate any noncompliant security group that is detected.
- C. In the transit account, create a VPC prefix list with all of the internal IP address ranges. Use AWS Resource Access Manager to share the prefix list with all of the other accounts. Use the shared prefix list to configure security group rules in the other accounts.
- D. In the transit account, create a security group with all of the internal IP address ranges. Configure the security groups in the other accounts to reference the transit account's security group by using a nested security group reference of `*<transit-account-id>./sg-1a2b3c4d`.

Answer: C

Explanation:

Customer-managed prefix lists — Sets of IP address ranges that you define and manage. You can share your prefix list with other AWS accounts, enabling those accounts to reference the prefix list in their own resources. <https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html> A VPC prefix list is created in the transit account with all of the internal IP address ranges, and then shared to all of

the other accounts using AWS Resource Access Manager. This allows for central management of the IP address ranges, and eliminates the need for manual updates to security group rules in each account. This solution also allows for compliance checks to be run using AWS Config and for any non-compliant security groups to be automatically remediated.

Question: 48

A company runs a new application as a static website in Amazon S3. The company has deployed the application to a production AWS account and uses Amazon CloudFront to deliver the website. The website calls an Amazon API Gateway REST API. An AWS Lambda function backs each API method.

The company wants to create a CSV report every 2 weeks to show each API Lambda function's recommended configured memory, recommended cost, and the price difference between current configurations and the recommendations. The company will store the reports in an S3 bucket.

Which solution will meet these requirements with the LEAST development time?

- A. Create a Lambda function that extracts metrics data for each API Lambda function from Amazon CloudWatch Logs for the 2-week period. Collate the data into tabular format. Store the data as a `_csvfile` in an S3 bucket. Create an Amazon Eventbridge rule to schedule the Lambda function to run every 2 weeks.
- B. Opt in to AWS Compute Optimizer. Create a Lambda function that calls the `ExportLambdaFunctionRecommendations` operation. Export the `_csv` file to an S3 bucket. Create an Amazon Eventbridge rule to schedule the Lambda function to run every 2 weeks.
- C. Opt in to AWS Compute Optimizer. Set up enhanced infrastructure metrics. Within the Compute Optimizer console, schedule a job to export the Lambda recommendations to a `_csvfile`. Store the file in an S3 bucket every 2 weeks.
- D. Purchase the AWS Business Support plan for the production account. Opt in to AWS Compute Optimizer for AWS Trusted Advisor checks. In the Trusted Advisor console, schedule a job to export the cost optimization checks to a `_csvfile`. Store the file in an S3 bucket every 2 weeks.

Answer: B

Explanation:

https://docs.aws.amazon.com/compute-optimizer/latest/APIReference/API_ExportLambdaFunctionRecommendations.html

Question: 49

The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.

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

- A. Activate the user-defined cost allocation tags that represent the application and the team.
- B. Activate the AWS generated cost allocation tags that represent the application and the team.
- C. Create a cost category for each application in Billing and Cost Management.
- D. Activate IAM access to Billing and Cost Management.
- E. Create a cost budget.
- F. Enable Cost Explorer.

Answer: A,C,F

Explanation:

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html><https://aws.amazon.com/premiumsupport/knowledge-center/cost-explorer-analyze-spending-and-usage/>

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html><https://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html>

The best combination of actions to meet the company's requirements is Options A, C, and F. Option A involves activating the user-defined cost allocation tags that represent the application and the team. This will allow the company to assign costs to different applications or teams, and will **allow them to be tracked in the monthly AWS bill.**

Option C involves creating a cost category for each application in Billing and Cost Management. This will allow the company to easily identify and compare costs across different applications and teams. Option F involves enabling Cost Explorer. This will allow the company to view the costs of their AWS resources over the last **12 months and to create forecasts for the next 12 months.**

These recommendations are in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that "You can use cost allocation tags to group your costs by application, team, or other categories" (Source:https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf). Additionally, the book states that "Cost Explorer enables you to view the costs of your AWS resources over the last 12 months and to create forecasts for the next 12 months" (Source:https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf).

Question: 50

An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in **each client's allow list.**

The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways **provide internet access to the private subnets.**

How should a solutions architect ensure that the web application can continue to call the third-party API after the migration?

- A. Associate a block of customer-owned public IP addresses to the VPC. Enable public IP addressing for public

subnets in the VPC.

B. Register a block of customer-owned public IP addresses in the AWS account. Create Elastic IP addresses from the address block and assign them to the NAT gateways in the VPC.

C. Create Elastic IP addresses from the block of customer-owned IP addresses. Assign the static Elastic IP addresses to the ALB.

D. Register a block of customer-owned public IP addresses in the AWS account. Set up AWS Global Accelerator to use Elastic IP addresses from the address block. Set the ALB as the accelerator

endpoint.

Answer: B

Explanation:

When EC2 instances reach third-party API through internet, their private IP addresses will be masked by NAT Gateway public IP address. <https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-bring-your-own-ip-byoip-for-amazon-vc/>

Question: 51

A company with several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
1
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAllActions",
      "Effect": "Allow", "Action": "*",
      "Resource":
    },
    {
      "Sid": "DenyCloudTrail", "Effect":
      "Deny",
      "Action": "cloudtrail:*" ,
      "Resource":
    }
  ]
}
```

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?

- A. Add s3:CreateBucket with Allow effect to the SCP.
- B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
- C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
- D. Remove the SCP from account 1111-1111-1111.

Answer: C

Explanation:

However A's explanation is incorrect -

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html "SCPs are similar to AWS Identity and Access Management (IAM) permission policies and use almost the same syntax. However, an SCP never grants permissions."

SCPs alone are not sufficient to granting permissions to the accounts in your organization. No permissions are granted by an SCP. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

Question: 52

A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.

Which solution will meet these requirements?

- A. Attach a role to the instance with permission to write to Amazon S3. Use the AWS Systems Manager Session Manager option to gain access to the instance and run commands to copy data into Amazon S3.
- B. Create an image of the instance with the reboot option turned on. Launch a new EC2 instance from the image. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.
- C. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.
- D. Create an image of the instance. Launch a new EC2 instance from the image. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

Answer: A

Explanation:

Taking a snapshot of the EBS volume using Amazon Data Lifecycle Manager (DLM) will meet the requirements because it allows you to create a backup of the volume without the need to access the instance or its SSH key pair. Additionally, DLM allows you to schedule the backups to occur at specific intervals and also enables you to copy the snapshots to an S3 bucket. This approach will not impact the running application as the backup is performed on the EBS volume level.

Question: 53

A solutions architect needs to copy data from an Amazon S3 bucket in 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 object ACLs 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:

Step B is necessary so that the user in the destination account has the necessary permissions to access the source bucket and list its contents, read its objects. Step D is needed so that the user in the destination account has the necessary permissions to access the destination bucket and list contents, put objects, and set object ACLs. Step F is necessary because the `aws s3 sync` command needs to be run using the IAM user credentials from the destination account, so that the objects will have the appropriate permissions for the user in the destination account once they are copied.

Question: 54

A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.

Which solution will meet these requirements?

- A. Create an alias for every new deployed version of the Lambda function. Use the AWS CLI `update-alias` command with the `routing-config` parameter to distribute the load.
- B. Deploy the application into a new CloudFormation stack. Use an Amazon Route 53 weighted routing policy to distribute the load.
- C. Create a version for every new deployed Lambda function. Use the AWS CLI `update-function-configuration` command with the `routing-config` parameter to distribute the load.
- D. Configure AWS CodeDeploy and use `CodeDeployDefault.OneAtATime` in the Deployment configuration to distribute the load.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-traffic-shifting/>
<https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html>

Question: 55

A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files are uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS `sftp.example.com` through the use of Amazon Route 53.

What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

- A. Move the EC2 instance into an Auto Scaling group. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record `sftp.example.com` in Route 53 to point to the ALB.
- B. Migrate the SFTP server to AWS Transfer for SFTP. Update the DNS record `sftp.example.com` in Route 53 to point to the server endpoint hostname.
- C. Migrate the SFTP server to a file gateway in AWS Storage Gateway. Update the DNS record `sftp.example.com` in Route 53 to point to the file gateway endpoint.
- D. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record `sftp.example.com` in Route 53 to point to the NLB.

Answer: B

Explanation: <https://aws.amazon.com/aws-transfer-family/faqs/>
<https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html>https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/?nc1=h_ls

Question: 56

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

What should the solutions architect do to meet these requirements?

- A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows File Server Use the SMB share to host the VMware data store. Use VM Import/Export to move the VMs to Amazon EC2.
- B. Use the VMware vSphere client to export the application as an image in Open Virealization Format (OVF) format Create an Amazon S3 bucket to store the image in the destination AWS Region. Create and apply an IAM role for VM Import Use the AWS CLI to run the EC2 import command.
- C. Configure AWS Storage Gateway for files service to export a Common Internet File System(CIFS) share. Create a backup copy to the shared folder. Sign in to the AWS Management Console and create an AMI from the backup copy Launch an EC2 instance that is based on the AMI.
- D. Create a managed-instance activation for a hybrid environment in AWS Systems Manager. Download and install Systems Manager Agent on the on-premises VM Register the VM with Systems Manager to be a managed instance Use AWS Backup to create a snapshot of the VM and create an AMI. Launch an EC2 instance that is based on the AMI

Answer: D

Explanation:

<https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html>

- Export an OVF Template
- Create / use an Amazon S3 bucket for storing the exported images. The bucket must be in the Region where you want to import your VMs.
- Create an IAM role named vmimport.
- You'll use AWS CLI to run the import commands.

<https://aws.amazon.com/premiumsupport/knowledge-center/import-instances/>

Question: 57

A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The

Lambda function is invoked when a new image is uploaded to Amazon S3.

The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.

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

- A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR).
- B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
- C. Create an AWS Step Functions state machine with a Parallel state to invoke the Lambda function. Increase the provisioned concurrency of the Lambda function.
- D. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of Amazon EC2. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
- E. Modify the application to store images on Amazon Elastic File System (Amazon EFS) and to store metadata on an Amazon RDS DB instance. Adjust the Lambda function to mount the EFS file share.

Answer: A,B

Explanation:

A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR). - This step is necessary to package the application code in a container and make it available for running on ECS. B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS taskdefinition when a new file arrives in Amazon S3.

Question: 58

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:

AWS Control Tower provides a set of "strongly recommended guardrails" that can be enabled to implement governance and policy enforcement. One of these guardrails is "Encrypt Amazon RDS instances" which will detect RDS DB instances that are not encrypted at rest. By enabling this guardrail and applying it to the production OU, the company will be able to enforce encryption for RDS instances in the production environment.

Question: 59

A startup company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.

The company's existing architecture includes the following:

- A VPC with private and public subnets, and a NAT gateway
- Site-to-Site VPN for connectivity with the on-premises environment
- EC2 security groups with direct SSH access from the on-premises environment

The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers.

Which strategy should a solutions architect use?

- A. Install and configure EC2 Instance Connect on the fleet of EC2 instances. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI.
- B. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's devices. Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.
- C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's devices. Enable AWS Config for EC2 security group resource changes. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.
- D. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attached. Attach the IAM role to all the EC2 instances. Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Answer: D

Explanation:

Allows client machines to be able to connect to Session Manager using the AWS CLI instead of going through the AWS EC2 or AWS Server Manager console. <https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.html><https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.html#:~:text=aws%20ssm%20start%2Dsession%20%2D%2Dtarget%20instance%2D>

Question: 60

A company that uses AWS Organizations allows developers to experiment on AWS. As part of the landing zone that the company has deployed, developers use their company email address to request an account. The company wants to ensure that developers are not launching costly services or running services unnecessarily. The company must give developers a fixed monthly budget to limit their AWS costs.

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

- A. Create an SCP to set a fixed monthly account usage limit. Apply the SCP to the developer accounts.
- B. Use AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process.
- C. Create an SCP to deny access to costly services and components. Apply the SCP to the developer accounts.
- D. Create an IAM policy to deny access to costly services and components. Apply the IAM policy to the developer accounts.
- E. Create an AWS Budgets alert action to terminate services when the budgeted amount is reached. Configure the action to terminate all services.
- F. Create an AWS Budgets alert action to send an Amazon Simple Notification Service (Amazon SNS) notification when the budgeted amount is reached. Invoke an AWS Lambda function to terminate all services.

Answer: B,C,F

Explanation:

Option A is incorrect because creating an SCP to set a fixed monthly account usage limit is not possible. SCPs are policies that specify the services and actions that users and roles can use in the member accounts of an AWS Organization. SCPs cannot enforce budget limits or prevent users from launching costly services or running services unnecessarily¹

Option B is correct because using AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets allows you to plan your service usage, service costs, and instance reservations. You can create budgets that alert you when your costs or usage exceed (or are forecasted to

exceed) your budgeted amount²

Option C is correct because creating an SCP to deny access to costly services and components meets the requirement of ensuring that developers are not launching costly services or running services unnecessarily. SCPs can restrict access to certain AWS services or actions based on conditions such as region, resource tags, or request time. For example, an SCP can deny access to Amazon Redshift clusters or Amazon EC2 instances with certain instance types¹

Option D is incorrect because creating an IAM policy to deny access to costly services and components is not sufficient to meet the requirement of ensuring that developers are not launching costly services or running services unnecessarily. IAM policies can only control access to resources within a single AWS account. If developers have multiple accounts or can create new accounts, they can bypass the IAM policy restrictions. SCPs can apply across multiple accounts within an AWS Organization and prevent users from creating new accounts that do not comply with the SCP rules³ Option E is incorrect because creating an AWS Budgets alert action to terminate services when the budgeted amount is reached is not possible. AWS Budgets alert actions can only perform one of the following actions: apply an IAM policy, apply an SCP, or send a notification through Amazon SNS. AWS Budgets alert actions cannot terminate services directly.

Option F is correct because creating an AWS Budgets alert action to send an Amazon SNS notification when the budgeted amount is reached and invoking an AWS Lambda function to terminate all services meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets alert actions can send notifications through Amazon SNS when a budget threshold is breached. Amazon SNS can trigger an AWS Lambda function that can perform custom logic such as terminating all services in the developer's account. This way, developers cannot exceed their budget limit and incur additional costs.

Reference: 1: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

2: <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets-create.html> 3:

<https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html> : <https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-actions.html> : <https://docs.aws.amazon.com/sns/latest/dg/sns-lambda.html> : <https://docs.aws.amazon.com/lambda/latest/dg/welcome.html>

Question: 61

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

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

Which solution will meet these requirements?

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

- C. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions and the Aurora DB cluster with the Target account. Grant the Target account permission to clone the Aurora DB cluster.
- D. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions with the Target account. Share the automated Aurora DB cluster snapshot with the Target account.

Answer: C

Explanation:

This solution uses a combination of AWS Resource Access Manager (RAM) and automated backups to migrate the Lambda functions and the Aurora database to the Target account while minimizing downtime. In this solution, the Lambda function deployment package is downloaded from the Source account and used to create new Lambda functions in the Target account. The Aurora DB cluster is shared with the Target account using AWS RAM and the Target account is granted permission to clone the Aurora DB cluster, allowing for a new copy of the Aurora database to be created in the Target account. This approach allows for the data to be migrated to the Target account while minimizing downtime, as the Target account can use the cloned Aurora database while the original Aurora database continues to be used in the Source account.

Question: 62

A company runs a Python script on an Amazon EC2 instance to process data.

a. The script runs every 10 minutes. The script ingests files from an Amazon S3 bucket and processes the files. On average, the script takes approximately 5 minutes to process each file. The script will not reprocess a file that the script has already processed.

The company reviewed Amazon CloudWatch metrics and noticed that the EC2 instance is idle for approximately 40% of the time because of the file processing speed. The company wants to make the workload highly available and scalable. The company also wants to reduce long-term management overhead.

Which solution will meet these requirements MOST cost-effectively?

- A. Migrate the data processing script to an AWS Lambda function. Use an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue. Configure Amazon S3 to send event notifications to the SQS queue. Create an EC2 Auto Scaling group with a minimum size of one instance. Update the data processing script to poll the SQS queue. Process the S3 objects that the SQS message identifies.
- C. Migrate the data processing script to a container image. Run the data processing container on an EC2 instance. Configure the container to poll the S3 bucket for new objects and to process the resulting objects.
- D. Migrate the data processing script to a container image that runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. Create an AWS Lambda function that calls the Fargate RunTaskAPI operation when the container processes the file. Use an S3 event notification to invoke the Lambda function.

Answer: A

Explanation:

migrating the data processing script to an AWS Lambda function and using an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects. This solution meets the company's requirements of high availability and scalability, as well as reducing long-term management overhead, and is likely to be the most cost-effective option.

Question: 63

A financial services company in North America plans to release a new online web application to its customers on AWS . The company will launch the application in the us-east-1 Region on Amazon EC2 instances. The application must be highly available and must dynamically scale to meet user traffic. The company also wants to implement a disaster recovery environment for the application in the us- west-1 Region by using active-passive failover.

Which solution will meet these requirements?

- A. Create a VPC in us-east-1 and a VPC in us-west-1 Configure VPC peering In the us-east-1VPC. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in both VPCs Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs Place the Auto Scaling group behind the ALB.
- B. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VPC. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPC. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC Place the Auto Scaling group behind the ALB Set up the same configuration in the us-west-1 VPC. Create an Amazon Route 53 hosted zone Create separate records for each ALB Enable health checks to ensure high availability between Regions.
- C. Create a VPC in us-east-1 and a VPC in us-west-1 In the us-east-1 VPC. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPC Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC Place the Auto Scaling group behind the ALB Set up the same configuration in the us-west-1 VPC Create an Amazon Route 53 hosted zone. Create separate records for each ALB Enable health checks and configure a failover routing policy for each record.
- D. Create a VPC in us-east-1 and a VPC in us-west-1 Configure VPC peering In the us-east-1 VPC. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs Place the Auto Scaling group behind the ALB Create an Amazon Route 53 host.. Create a record for the ALB.

Answer: C

Explanation:

it's the one that handles failover while B (the one shown as the answer today) it almost the same but does not handle failover.

Question: 64

A company has an environment that has a single AWS account. A solutions architect is reviewing the environment to recommend what the company could improve specifically in terms of access to the AWS Management Console. The company's IT support workers currently access the console for administrative tasks, authenticating with named IAM users that have been mapped to their job role. The IT support workers no longer want to maintain both their Active Directory and IAM user accounts. They want to be able to access the console by using their existing Active Directory credentials. The solutions architect is using AWS Single Sign-On (AWS SSO) to implement this functionality.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an organization in AWS Organizations. Turn on the AWS SSO feature in Organizations Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Directory. Configure AWS SSO and set the AWS Managed Microsoft AD directory as the identity source. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.
- B. Create an organization in AWS Organizations. Turn on the AWS SSO feature in Organizations Create and configure an AD Connector to connect to the company's on-premises Active Directory. Configure AWS SSO and select the AD Connector as the identity source. Create permission sets and map them to the existing groups within the company's Active Directory.
- C. Create an organization in AWS Organizations. Turn on all features for the organization. Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Directory. Configure AWS SSO and select the AWS Managed Microsoft AD directory as the identity source. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.
- D. Create an organization in AWS Organizations. Turn on all features for the organization. Create and configure an AD Connector to connect to the company's on-premises Active Directory. Configure AWS SSO and select the AD Connector as the identity source. Create permission sets and map them to the existing groups within the company's Active Directory.

Answer: D

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_org_support-all-features.html
<https://docs.aws.amazon.com/singlesignon/latest/userguide/get-started-prereqs-considerations.html>

Question: 65

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

for these uploads

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

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

Answer: A,D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/> Enabling S3 Transfer Acceleration on the S3 bucket and configuring the app to use the Transfer Acceleration endpoint for uploads will improve the app's performance for these uploads by leveraging Amazon CloudFront's globally distributed edge locations to accelerate the uploads. Breaking the video files into chunks and using a multipart upload to transfer files to Amazon S3 will also improve the app's performance by allowing parts of the file to be uploaded in parallel, reducing the overall upload time.

Question: 66

An application is using an Amazon RDS for MySQL Multi-AZ DB instance in the us-east-1 Region. After a failover test, the application lost the connections to the database and could not re-establish the connections. After a restart of the application, the application re-established the connections.

A solutions architect must implement a solution so that the application can re-establish connections to the database without requiring a restart.

Which solution will meet these requirements?

- A. Create an Amazon Aurora MySQL Serverless v1 DB instance. Migrate the RDS DB instance to the Aurora Serverless v1 DB instance. Update the connection settings in the application to point to the Aurora reader endpoint.
- B. Create an RDS proxy. Configure the existing RDS endpoint as a target. Update the connection settings in the application to point to the RDS proxy endpoint.
- C. Create a two-node Amazon Aurora MySQL DB cluster. Migrate the RDS DB instance to the Aurora DB cluster. Create an RDS proxy. Configure the existing RDS endpoint as a target. Update the connection settings in the application to point to the RDS proxy endpoint.
- D. Create an Amazon S3 bucket. Export the database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Configure Amazon Athena to use the S3 bucket as a data store. Install the latest Open Database Connectivity (ODBC) driver for the application. Update the connection settings in the application to point to

the Athena endpoint

Answer: B

Explanation:

Amazon RDS Proxy is a fully managed database proxy service for Amazon Relational Database Service (RDS) that makes applications more scalable, resilient, and secure. It allows applications to pool and share connections to an RDS database, which can help reduce database connection overhead, improve scalability, and provide automatic failover and high availability.

Question: 67

A company is building a solution in the AWS Cloud. Thousands of devices will connect to the solution and send data.

a. Each device needs to be able to send and receive data in real time over the MQTT protocol. Each device must authenticate by using a unique X.509 certificate.

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

- A. Set up AWS IoT Core. For each device, create a corresponding Amazon MQ queue and provision a certificate. Connect each device to Amazon MQ.
- B. Create a Network Load Balancer (NLB) and configure it with an AWS Lambda authorizer. Run an MQTT broker on Amazon EC2 instances in an Auto Scaling group. Set the Auto Scaling group as the target for the NLB. Connect each device to the NLB.
- C. Set up AWS IoT Core. For each device, create a corresponding AWS IoT thing and provision a certificate. Connect each device to AWS IoT Core.
- D. Set up an Amazon API Gateway HTTP API and a Network Load Balancer (NLB). Create integration between API Gateway and the NLB. Configure a mutual TLS certificate authorizer on the HTTP API. Run an MQTT broker on an Amazon EC2 instance that the NLB targets. Connect each device to the NLB.

Answer: C

Explanation:

This solution requires minimal operational overhead, as it only requires setting up AWS IoT Core and creating a thing for each device. (Reference: AWS Certified Solutions Architect - Professional Official Amazon Text

Book, Page 537)

AWS IoT Core is a fully managed service that enables secure, bi-directional communication between internet-connected devices and the AWS Cloud. It supports the MQTT protocol and includes built-in device authentication and access control. By using AWS IoT Core, the company can easily provision and manage the X.509 certificates for each device, and connect the devices to the service with minimal operational overhead.

Question: 68

A company is running several workloads in a single AWS account. A new company policy states that engineers

can provision only approved resources and that engineers must use AWS CloudFormation to provision these resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.

What should the solutions architect do to create the solution?

- A. Upload AWS CloudFormation templates that contain approved resources to an Amazon S3 bucket. Update the IAM policy for the engineers' IAM role to only allow access to Amazon S3 and AWS CloudFormation. Use AWS CloudFormation templates to provision resources.
- B. Update the IAM policy for the engineers' IAM role with permissions to only allow provisioning of approved resources and AWS CloudFormation. Use AWS CloudFormation templates to create stacks with approved resources.
- C. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation actions. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service role. Assign the IAM service role to AWS CloudFormation during stack creation.
- D. Provision resources in AWS CloudFormation stacks. Update the IAM policy for the engineers' IAM role to only allow access to their own AWS CloudFormation stack.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/security-best-practices.html#use-iam-to-control-access><https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-service-role.html>

Question: 69

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:

DynamoDB with TTL, cheaper for sustained throughput of small items + suited for fast retrievals. S3 cheaper for storage only, much higher costs with writes. RDS not designed for this use case.

Question: 70

A retail company is hosting an ecommerce website on AWS across multiple AWS Regions. The company wants the website to be operational at all times for online purchases. The website stores data in an Amazon RDS for MySQL DB instance.

Which solution will provide the HIGHEST availability for the database?

- A. Configure automated backups on Amazon RDS. In the case of disruption, promote an automated backup to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.
- B. Configure global tables and read replicas on Amazon RDS. Activate the cross-Region scope. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.
- C. Configure global tables and automated backups on Amazon RDS. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.
- D. Configure read replicas on Amazon RDS. In the case of disruption, promote a cross-Region and read replica to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.

Answer: D

Explanation:

This solution will provide the highest availability for the database, as the read replicas will allow the database to be available in multiple Regions, thus reducing the chances of disruption. Additionally, the promotion of the cross-Region read replica to become a standalone DB instance will ensure that the database is still available even if one of the Regions experiences disruptions.

Question: 71

A company recently completed the migration from an on-premises data center to the AWS Cloud by using a replatforming strategy. One of the migrated servers is running a legacy Simple Mail Transfer Protocol (SMTP) service that a critical application relies upon. The application sends outbound email messages to the company's customers. The legacy SMTP server does not support TLS encryption and uses TCP port 25. The application can use SMTP only.

The company decides to use Amazon Simple Email Service (Amazon SES) and to decommission the legacy SMTP server. The company has created and validated the SES domain. The company has lifted the SES limits.

What should the company do to modify the application to send email messages from Amazon SES?

- A. Configure the application to connect to Amazon SES by using TLS Wrapper. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Attach the IAM role to an Amazon EC2 instance.
- B. Configure the application to connect to Amazon SES by using STARTTLS. Obtain Amazon SES SMTP credentials. Use the credentials to authenticate with Amazon SES.
- C. Configure the application to use the SES API to send email messages. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Use the IAM role as a service role for Amazon SES.
- D. Configure the application to use AWS SDKs to send email messages. Create an IAM user for Amazon SES. Generate API access keys. Use the access keys to authenticate with Amazon SES.

Answer: B

Explanation:

To set up a STARTTLS connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 25, 587, or 2587, issues an EHLO command, and waits for the server to announce that it supports the STARTTLS SMTP extension. The client then issues the STARTTLS command, initiating TLS negotiation. When negotiation is complete, the client issues an EHLO command over the new encrypted connection, and the SMTP session proceeds normally. To set up a TLS Wrapper connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 465 or 2465. The server presents its certificate, the client issues an EHLO command, and the SMTP session proceeds normally.

<https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html>

Question: 72

A company recently acquired several other companies. Each company has a separate AWS account with a different billing and reporting method. The acquiring company has consolidated all the accounts into one organization in AWS Organizations. However, the acquiring company has found it difficult to generate a cost report that contains meaningful groups for all the teams.

The acquiring company's finance team needs a solution to report on costs for all the companies through a self-managed application.

Which solution will meet these requirements?

- A. Create an AWS Cost and Usage Report for the organization. Define tags and cost categories in the report. Create a table in Amazon Athena. Create an Amazon QuickSight dataset based on the Athena table. Share the dataset with the finance team.
- B. Create an AWS Cost and Usage Report for the organization. Define tags and cost categories in the report. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.

- C. Create an Amazon QuickSight dataset that receives spending information from the AWS Price List Query API. Share the dataset with the finance team.
- D. Use the AWS Price List Query API to collect account spending information. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.

Answer: A

Explanation:

Creating an AWS Cost and Usage Report for the organization and defining tags and cost categories in the report will allow for detailed cost reporting for the different companies that have been consolidated into one organization. By creating a table in Amazon Athena and an Amazon QuickSight dataset based on the Athena table, the finance team will be able to easily query and generate reports on the costs for all the companies. The dataset can then be shared with the finance team for them to use for their reporting needs.

Question: 73

A company runs an IoT platform on AWS IoT 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 and add 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: B,E

Explanation:

Option C is correct because leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data resolves the issues permanently and enable growth as new sensors are provisioned. Amazon Kinesis Data Streams is a serverless streaming data service that simplifies the capture, processing, and storage of data streams at any scale. Kinesis Data Streams can handle any amount of streaming data and process data from hundreds of thousands of sources with very low latency. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be triggered by Kinesis Data Streams events and process the data records in real time. Lambda can also scale automatically based on the incoming data volume. By using Kinesis Data Streams and Lambda, the company can reduce the load on the API servers and improve the performance and scalability of the data ingestion and processing layer.

Option E is correct because re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance resolves the issues permanently and enable growth as new sensors are provisioned. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB supports auto scaling, which automatically adjusts read and write capacity based on actual traffic patterns. DynamoDB also supports on-demand capacity mode, which instantly accommodates up to double the previous peak traffic on a table. By using DynamoDB instead of RDS MySQL DB instance, the company can eliminate high write latency and improve scalability and performance of the database tier.

Reference: 1: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html> 2: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html 3: <https://docs.aws.amazon.com/streams/latest/dev/introduction.html> : <https://docs.aws.amazon.com/lambda/latest/dg/welcome.html> : <https://docs.aws.amazon.com/xray/latest/devguide/aws-xray.html> : <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html> :

Question: 74

A company is building an electronic document management system in which users upload their documents. The application stack is entirely serverless and runs on AWS in the eu-central-1 Region. The system includes a web application that uses an Amazon CloudFront distribution for delivery with Amazon S3 as the origin. The web application communicates with Amazon API Gateway Regional endpoints. The API Gateway APIs call AWS Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket.

The company is growing steadily and has completed a proof of concept with its largest customer. The company must improve latency outside of Europe.

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

- A. Enable S3 Transfer Acceleration on the S3 bucket. Ensure that the web application uses the Transfer Acceleration signed URLs.
- B. Create an accelerator in AWS Global Accelerator. Attach the accelerator to the CloudFront distribution.
- C. Change the API Gateway Regional endpoints to edge-optimized endpoints.
- D. Provision the entire stack in two other locations that are spread across the world. Use global databases on the Aurora Serverless cluster.
- E. Add an Amazon RDS proxy between the Lambda functions and the Aurora Serverless database.

Answer: A,C

Explanation:

<https://aws.amazon.com/global-accelerator/faqs/>

Question: 75

An adventure company has launched a new feature on its mobile app. Users can use the feature to upload their hiking and ratting photos and videos anytime. The photos and videos are stored in Amazon S3 Standard storage

in an S3 bucket and are served through Amazon CloudFront.

The company needs to optimize the cost of the storage. A solutions architect discovers that most of the uploaded photos and videos are accessed infrequently after 30 days. However, some of the uploaded photos and videos are accessed frequently after 30 days. The solutions architect needs to implement a solution that maintains millisecond retrieval availability of the photos and videos at the lowest possible cost.

Which solution will meet these requirements?

- A. Configure S3 Intelligent-Tiering on the S3 bucket.
- B. Configure an S3 Lifecycle policy to transition image objects and video objects from S3 Standard to S3 Glacier Deep Archive after 30 days.
- C. Replace Amazon S3 with an Amazon Elastic File System (Amazon EFS) file system that is mounted on Amazon EC2 instances.
- D. Add a Cache-Control: max-age header to the S3 image objects and S3 video objects. Set the header to 30 days.

Answer: A

Explanation:

Amazon S3 Intelligent-Tiering is a storage class that automatically moves objects between two access tiers based on changing access patterns. Objects that are accessed frequently are stored in the frequent access tier and objects that are accessed infrequently are stored in the infrequent access tier. This allows for cost optimization without requiring manual intervention. This makes it an ideal solution for the scenario described, as it can automatically move objects that are infrequently accessed after 30 days to a lower-cost storage tier while still maintaining millisecond retrieval availability.

Question: 76

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 identify 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: B

Explanation:

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

Question: 77

A company has its cloud infrastructure on AWS. A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts. What should the solutions architect do to meet these requirements?

- A. Use AWS CloudFormation templates. Add IAM policies to control the various accounts. Deploy the templates across the multiple Regions.
- B. Use AWS Organizations. Deploy AWS CloudFormation templates from the management account. Use AWS Control Tower to manage deployments across accounts.
- C. Use AWS Organizations and AWS CloudFormation StackSets. Deploy a CloudFormation template from an account that has the necessary IAM permissions.
- D. Use nested stacks with AWS CloudFormation templates. Change the Region by using nested stacks.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-organization/>

AWS Organizations allows the management of multiple AWS accounts as a single entity and AWS CloudFormation StackSets allows creating, updating, and deleting stacks across multiple accounts and regions in an organization. This solution allows creating a single CloudFormation template that can be deployed across multiple accounts and regions, and also allows for the management of access and permissions for the different accounts through the use of IAM roles and policies in the management account.

Question: 78

A company plans to refactor a monolithic application into a modern application designed for AWS. The CI/CD pipeline needs to be upgraded to support the modern design for the application with the following requirements:

- It should allow changes to be released several times every hour.
 - * It should be able to roll back the changes as quickly as possible.
- Which design will meet these requirements?

- A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurations. Deploy the application by replacing Amazon EC2 instances.
- B. Specify AWS Elastic Beanstalk to stage in a secondary environment as the deployment target for the CI/CD pipeline of the application. To deploy, swap the staging and production environment URLs.

- C. Use AWS Systems Manager to re-provision the infrastructure for each deployment Update the Amazon EC2 user data to pull the latest code art-fact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment
- D. Roll out the application updates as pan of an Auto Scaling event using prebuilt AMIs. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

Explanation:

It is the fastest when it comes to rollback and deploying changes every hour

Question: 79

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

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

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

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

Answer: A,B

Explanation:

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. This allows the instances to make outbound connections to the DB cluster on 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. This allows connections to the DB cluster from the EC2 instances on the default Aurora port.

Question: 80

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

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

- A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in each account to create monthly reports for each business unit.
- B. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
- C. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.
- D. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owner. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

Explanation:

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

<https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Budgets%20gives%20you%20the,below%20the%20threshold%20you%20define>

Question: 81

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

- ManagementConsole. 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 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:

This solution will meet the requirements of the company as it provides a secure, cost-effective and fast way of transferring large data sets from on-premises to AWS. Snowball Edge devices encrypt the data during transfer, and the devices are shipped back to AWS for import into S3. This option is more cost effective than using Direct Connect or VPN connections as it does not require the company to pay for long-term dedicated connections.

Question: 82

A company has migrated its forms-processing application to AWS. When users interact with the application, they upload scanned forms as files through a web application. A database stores user metadata and references to files that are stored in Amazon S3. The web application runs on Amazon EC2 instances and an Amazon RDS for PostgreSQL database.

When forms are uploaded, the application sends notifications to a team through Amazon Simple Notification Service (Amazon SNS). A team member then logs in and processes each form. The team member performs data validation on the form and extracts relevant data before entering the information into another system that uses an API.

A solutions architect needs to automate the manual processing of the forms. The solution must provide accurate form extraction, minimize time to market, and minimize long-term operational overhead. Which solution will meet these requirements?

- A. Develop custom libraries to perform optical character recognition (OCR) on the forms. Deploy the libraries to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster as an application tier. Use this tier to process the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data into an Amazon DynamoDB table. Submit the data to the target system's API. Host the new application tier on EC2 instances.
- B. Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use artificial intelligence and machine learning (AI/ML) models that are trained and hosted on an EC2 instance to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API.
- C. Host a new application tier on EC2 instances. Use this tier to call endpoints that host artificial intelligence and machine learning (AI/ML) models that are trained and hosted in Amazon SageMaker to perform optical character recognition (OCR) on the forms. Store the output in Amazon ElastiCache. Parse this output by

extracting the data that is required within the application tier. Submit the data to the target system's API.

D. Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API.

Answer: D

Explanation:

Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's

API. This solution meets the requirements of accurate form extraction, minimal time to market, and minimal long-term operational overhead. Amazon Textract and Amazon Comprehend are fully managed and serverless services that can perform OCR and extract relevant data from the forms, which eliminates the need to develop custom libraries or train and host models. Using AWS Step Functions and Lambda allows for easy automation of the process and the ability to scale as needed.

Question: 83

A company is refactoring its on-premises order-processing platform in the AWS Cloud. The platform includes a web front end that is hosted on a fleet of VMs RabbitMQ to connect the front end to the backend, and a Kubernetes cluster to run a containerized backend system to process the orders. The company does not want to make any major changes to the application. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an AMI of the web server VM. Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer. Set up Amazon MQ to replace the on-premises messaging queue. Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend.
- B. Create a custom AWS Lambda runtime to mimic the web server environment. Create an Amazon API Gateway API to replace the front-end web servers. Set up Amazon MQ to replace the on-premises messaging queue. Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend.
- C. Create an AMI of the web server VM. Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer. Set up Amazon MQ to replace the on-premises messaging queue. Install Kubernetes on a fleet of different EC2 instances to host the order-processing backend.
- D. Create an AMI of the web server VM. Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer. Set up an Amazon Simple Queue Service (Amazon SQS) queue to replace the on-premises messaging queue. Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend.

Answer: A

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2020/11/announcing-amazon-mq-rabbitmq/>

Question: 84

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:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "kms:GenerateDataKey",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "kms:Decrypt",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "kms:DescribeKey",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "kms:Encrypt",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "kms:RevokeDataKey",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "kms:UpdateKey",
      "Resource": "*"
    }
  ]
}
```

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

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

- A. `Kms:GenerateDataKey`
- B. `Kms:GetKeyPolicy`
- C. `kms:GetPublicKey`

D. kms:SKjn

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-access-denied-error-kms/>

"An error occurred (AccessDenied) when calling the PutObject operation: Access Denied" This error message indicates that your IAM user or role needs permission for the kms:GenerateDataKey action.

Question: 85 A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.

How should a solutions architect configure the web ACLs to meet these requirements?

- A. Set the action of the web ACL rules to Count. Enable AWS WAF logging Analyze the requests for false positives Modify the rules to avoid any false positive Over time change the action of the web ACL rules from Count to Block.
- B. Use only rate-based rules in the web ACLs. and set the throttle limit as high as possible Temporarily block all requests that exceed the limit. Define nested rules to narrow the scope of the rate tracking.
- C. Set the action of the web ACL rules to Block. Use only AWS managed rule groups in the web ACLs Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.
- D. Use only custom rule groups in the web ACLs. and set the action to Allow Enable AWS WAF logging Analyze the requests for false positives Modify the rules to avoid any false positive Over time, change the action of the web ACL rules from Allow to Block.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/waf-analyze-count-action-rules/>

Question: 86

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 solution meets the requirements with the least amount of operational overhead as it requires the security team to create and maintain a single customer-managed prefix list, and share it with the organization using AWS Resource Access Manager. The owners of each AWS account are then responsible for allowing the prefix list in their security groups, which eliminates the need for the security team to manually notify each account owner when changes are made. This solution also eliminates the need for a separate AWS Lambda function in each account, reducing the overall complexity of the solution.

Question: 87

A company has introduced a new policy that allows employees to work remotely from their homes if they connect by using a VPN. The company is hosting internal applications with VPCs in multiple AWS accounts. Currently, the applications are accessible from the company's on-premises office network through an AWS Site-to-Site VPN connection. The VPC in the company's main AWS account has peering connections established with VPCs in other AWS accounts.

A solutions architect must design a scalable AWS Client VPN solution for employees to use while they work from home.

What is the MOST cost-effective solution that meets these requirements?

- A. Create a Client VPN endpoint in each AWS account Configure required routing that allows access to internal applications
- B. Create a Client VPN endpoint in the main AWS account Configure required routing that allows access to internal applications
- C. Create a Client VPN endpoint in the main AWS account Provision a transit gateway that is connected to each AWS account Configure required routing that allows access to internal applications
- D. Create a Client VPN endpoint in the main AWS account Establish connectivity between the Client VPN endpoint and the AWS Site-to-Site VPN

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/scenario-peered.html>

Question: 88

A company is running an application in the AWS Cloud. Recent application metrics show inconsistent response times and a significant increase in error rates. Calls to third-party services are causing the delays. Currently, the application calls third-party services synchronously by directly invoking an AWS Lambda function.

A solutions architect needs to decouple the third-party service calls and ensure that all the calls are eventually completed.

Which solution will meet these requirements?

- A. Use an Amazon Simple Queue Service (Amazon SQS) queue to store events and invoke the Lambda function.
- B. Use an AWS Step Functions state machine to pass events to the Lambda function.
- C. Use an Amazon EventBridge rule to pass events to the Lambda function.
- D. Use an Amazon Simple Notification Service (Amazon SNS) topic to store events and Invoke the Lambda function.

Answer: A

Explanation:

Using an SQS queue to store events and invoke the Lambda function will decouple the third-party service calls and ensure that all the calls are eventually completed. SQS allows you to store messages in a queue and process them asynchronously, which eliminates the need for the application to wait for a response from the third-party service. The messages will be stored in the SQS queue until they are processed by the Lambda function, even if the Lambda function is currently unavailable or busy. This will ensure that all the calls are eventually completed, even if there are delays or errors. AWS Step Functions state machines can also be used to pass events to the

Lambda function, but it would require additional management and configuration to set up the state machine, which would increase operational overhead.

Amazon EventBridge rule can also be used to pass events to the Lambda function, but it would not provide the same level of decoupling and reliability as SQS.

Using Amazon Simple Notification Service (Amazon SNS) topic to store events and Invoke the Lambda function, is similar to SQS, but SNS is a publish-subscribe messaging service and SQS is a queue service. SNS is used for sending messages to multiple recipients, SQS is used for sending messages to a single recipient, so SQS is more appropriate for this use case.

Reference:

AWS SQS

AWS Step Functions

AWS EventBridge

AWS SNS

Question: 89

A company is running applications on AWS in a multi-account environment. The company's sales team and marketing team use separate AWS accounts in AWS Organizations.

The sales team stores petabytes of data in an Amazon S3 bucket. The marketing team uses Amazon QuickSight for data visualizations. The marketing team needs access to data that the sales team stores in the S3 bucket. The company has encrypted the S3 bucket with an AWS Key Management Service (AWS KMS) key. The marketing team has already created the IAM service role for QuickSight to provide QuickSight access in the marketing AWS account. The company needs a solution that will provide secure access to the data in the S3 bucket across AWS accounts.

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

- A. Create a new S3 bucket in the marketing account. Create an S3 replication rule in the sales account to copy the objects to the new S3 bucket in the marketing account. Update the QuickSight permissions in the marketing account to grant access to the new S3 bucket.
- B. Create an SCP to grant access to the S3 bucket to the marketing account. Use AWS Resource Access Manager (AWS RAM) to share the KMS key from the sales account with the marketing account. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
- C. Update the S3 bucket policy in the marketing account to grant access to the QuickSight role. Create a KMS grant for the encryption key that is used in the S3 bucket. Grant decrypt access to the QuickSight role. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
- D. Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight role, to create a trust relationship with the new IAM role in the sales account.

Answer: D

Explanation:

Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight role, to create a trust relationship with the new IAM role in the sales account.

This approach is the most secure way to grant cross-account access to the data in the S3 bucket while minimizing operational overhead. By creating an IAM role in the sales account, the marketing team can assume the role in their own account, and have access to the S3 bucket. And updating the

QuickSight role, to create a trust relationship with the new IAM role in the sales account will grant the marketing team to access the data in the S3 bucket and use it for data visualization using QuickSight.

AWS Resource Access Manager (AWS RAM) also allows sharing of resources between accounts, but it would require additional management and configuration to set up the sharing, which would increase operational overhead.

Using S3 replication would also replicate the data to the marketing account, but it would not provide the marketing team access to the original data, and also it would increase operational overhead with managing the replication process.

IAM roles and policies, KMS grants and trust relationships are a powerful combination for managing cross-account access in a secure and efficient manner.

Reference:

AWS IAM Roles AWS KMS - Key Grants AWS RAM

Question: 90

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

Which solution will meet these requirements?

- A. Migrate the SQL Server databases to Amazon RDS for MySQL by using backup and restore utilities.
- B. Use an AWS Snowball Edge Storage Optimized device to transfer data to Amazon S3. Set up Amazon RDS for MySQL. Use S3 integration with SQL Server features, such as BULK INSERT.
- C. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MeSQL. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.
- D. Use AWS DataSync to migrate data over the network between on-premises storage and Amazon S3. Set up Amazon RDS for MySQL. Use S3 integration with SQL Server features, such as BULK INSERT.

Answer: C

Explanation:

<https://aws.amazon.com/dms/schema-conversion-tool/>

AWS Schema Conversion Tool (SCT) can automatically convert the database schema from Microsoft SQL Server to Amazon RDS for MySQL. This allows for a smooth transition of the database schema without any manual intervention. AWS DMS can then be used to migrate the data from the onpremises databases to the newly

created Amazon RDS for MySQL instance. This service can perform a one-time migration of the data or can set up ongoing replication of data changes to keep the on-premises and AWS databases in sync.

Question: 91

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

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

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

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

Answer: A,C,E

Explanation:

A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket. The policy grants the necessary permissions to access the assets in the production S3 bucket. C

. In the production account, create a role. Attach the new policy to the role. Define the development account as a trusted entity. By creating a role and attaching the policy, and then defining the development account as a trusted entity, the development account can assume the role and access the production S3 bucket with the read and write permissions.

E. In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account. The IAM policy attached to the group allows the design team members to assume the role created in the production account, thereby giving them access to the production S3 bucket.

Step 1: Create a role in the Production Account; create the role in the Production account and specify the Development account as a trusted entity. You also limit the role permissions to only read and write access to the productionapp bucket. Anyone granted permission to use the role can read and write to the productionapp bucket. Step 2: Grant access to the role Sign in as an administrator in the Development account and allow the AssumeRole action on the UpdateApp role in the Production account. So, recap, production account you create

the policy for S3, and you set development account as a trusted entity. Then on the development account you allow the sts:assumeRole action on the role in production account.https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

Question: 92

A company developed a pilot application by using AWS Elastic Beanstalk and Java

a. To save costs during development, the company's development team deployed the application into a single-instance environment. Recent tests indicate that the application consumes more CPU than expected. CPU utilization is regularly greater than 85%, which causes some performance **bottlenecks**.

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

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

- A. Create a new Elastic Beanstalk application. Select a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the maximum CPU utilization is over 85% for 5 minutes.
- B. Create a second Elastic Beanstalk environment. Apply the traffic-splitting deployment policy. Specify a percentage of incoming traffic to direct to the new environment in the average CPU utilization is over 85% for 5 minutes.
- C. Modify the existing environment's capacity configuration to use a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.
- D. Select the Rebuild environment action with the load balancing option. Select an Availability Zones. Add a scale-out rule that will run if the sum CPU utilization is over 85% for 5 minutes.

Answer: C

Explanation:

This solution will meet the requirements with the least operational overhead because it allows the company to modify the existing environment's capacity configuration, so it becomes a load-balanced environment type. By selecting all availability zones, the company can ensure that the application is running in multiple availability zones, which can help to improve the availability and scalability of the application. The company can also add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes, which can help to mitigate the performance issues. This solution does not require creating new Elastic Beanstalk environments or rebuilding the existing one, which reduces the operational overhead.

You can refer to the AWS Elastic Beanstalk documentation for more information on how to use this service:<https://aws.amazon.com/elasticbeanstalk/> You can refer to the AWS documentation for more information on how to use autoscaling:<https://aws.amazon.com/autoscaling/>

Question: 93

A finance company is running its business-critical application on current-generation Linux EC2 instances. The

application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.

Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

- A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.
- B. Performing a one-time migration of the database cluster to Amazon RDS, and creating several additional read replicas to handle the load during end of month
- C. Using Amazon CloudWatch with AWS Lambda to change the type, size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric
- D. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Answer: B

Explanation:

In this scenario, the Amazon EC2 instances are in an Auto Scaling group already which means that the database read operations is the possible bottleneck especially during the month-end wherein the reports are generated. This can be solved by creating RDS read replicas.

Question: 94

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:

According to the AWS documentation¹, AWS App2Container (A2C) is a command line tool for migrating and modernizing Java and .NET web applications into container format. AWS A2C analyzes and builds an inventory of applications running in bare metal, virtual machines, Amazon Elastic Compute Cloud (EC2) instances, or in the cloud. You can use AWS A2C to generate container images for your applications and deploy them on Amazon ECS or Amazon EKS.

Option A meets the requirements of the scenario because it allows you to migrate your existing Java application to AWS and minimize the administrative overhead to maintain the servers. You can use AWS A2C to analyze your application dependencies, extract application artifacts, and generate a Dockerfile. You can then store your container images in Amazon ECR, which is a fully managed container registry service. You can use AWS Fargate as the launch type for your Amazon ECS cluster, which is a serverless compute engine that eliminates the need to provision and manage servers for your containers. You can grant the ECS task execution role permission to access the ECR image repository, which allows your tasks to pull images from ECR. You can configure Amazon ECS to use an ALB, which is a load balancer that distributes traffic across multiple targets in multiple Availability Zones using HTTP or HTTPS protocols. You can use the ALB to interact with your application.

Question: 95

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 to 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: B

Explanation:

This solution allows for deploying the Lambda function and API Gateway endpoint to another region, providing a failover option in case of any issues in the primary region. Using Route 53's failover routing policy allows for automatic routing of traffic to the healthy endpoint, ensuring that the application is available even in case of issues in one region. This solution provides a cost-effective and simple way to implement failover while minimizing operational overhead.

Question: 96

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 R1 sharing.
- B. Remove the HR department's production AWS account from the organization. Add the account to the consolidating billing configuration only.
- C. In the AWS Billing and Cost Management console, use the organization's management account to turn off R1 sharing for the HR department's 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:

You can use the management account of the organization in AWS Billing and Cost Management console to turn off RI sharing for the HR department's production AWS account. This will prevent other departments from sharing the RI discounts and ensure that only the HR department can use the RIs purchased in their production account.

Question: 97

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

Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive

How should the solutions architect gain access to an EC2 instance to troubleshoot the issue?

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

Answer: D

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html>

Question: 98

A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts.

The accounts are managed under different OUs in AWS Organizations.

Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed.

Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts.

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

- A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organization. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage. Update the parameter as needed to add or remove accounts or OUs. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account.
- B. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rules. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resources. Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
- C. Create AWS WAF rules in the management account of the organization. Use AWS Lambda environment variables to store account numbers and OUs to manage. Update environment variables as needed to add or remove accounts or OUs. Create cross-account IAM roles in member accounts. Assume the roles by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts.
- D. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization. Use AWS

Key Management Service (AWS KMS) to store account numbers and OUs to manage. Update AWS KMS as needed to add or remove accounts or OUs. Create IAM users in member accounts. Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules.

in the member accounts

Answer: A

Explanation:

<https://aws.amazon.com/solutions/implementations/automations-for-aws-firewall-manager/>

In this solution, AWS Firewall Manager is used to manage AWS WAF rules across accounts in the organization. An AWS Systems Manager Parameter Store parameter is used to store account numbers and OUs to manage. This parameter can be updated as needed to add or remove accounts or OUs. An Amazon EventBridge rule is used to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account. This solution allows for easy management of AWS WAF rules across multiple accounts with minimal operational overhead

Question: 99

A solutions architect is auditing the security setup of an AWS Lambda function for a company. The Lambda function retrieves the latest changes from an Amazon Aurora database. The Lambda function and the database run in the same VPC. Lambda environment variables are providing the database credentials to the Lambda function.

The Lambda function aggregates data and makes the data available in an Amazon S3 bucket that is configured for server-side encryption with AWS KMS managed encryption keys (SSE-KMS). The data must not travel across the internet. If any database credentials become compromised, the company needs a solution that minimizes the impact of the compromise.

What should the solutions architect recommend to meet these requirements?

- A. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
- B. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Enforce HTTPS on the connection to Amazon S3 during data transfers.
- C. Save the database credentials in AWS Systems Manager Parameter Store. Set up password rotation on the credentials in Parameter Store. Change the IAM role for the Lambda function to allow the function to access Parameter Store. Modify the Lambda function to retrieve the credentials from Parameter Store. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
- D. Save the database credentials in AWS Secrets Manager. Set up password rotation on the credentials in Secrets Manager. Change the IAM role for the Lambda function to allow the function to access Secrets Manager. Modify the Lambda function to retrieve the credentials from Secrets Manager. Enforce HTTPS on the connection to Amazon S3 during data transfers.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.IAMDBAuth.htm>

Question: 100

A large mobile gaming company has successfully migrated all of its on-premises infrastructure to the AWS Cloud. A solutions architect is reviewing the environment to ensure that it was built according to the design and that it is running in alignment with the Well-Architected Framework.

While reviewing previous monthly costs in Cost Explorer, the solutions architect notices that the creation and subsequent termination of several large instance types account for a high proportion of the costs. The solutions architect finds out that the company's developers are launching new Amazon EC2 instances as part of their testing and that the developers are not using the appropriate instance types.

The solutions architect must implement a control mechanism to limit the instance types that only the developers can launch.

Which solution will meet these requirements?

- A. Create a desired-instance-type managed rule in AWS Config. Configure the rule with the instance types that are allowed. Attach the rule to an event to run each time a new EC2 instance is launched.
- B. In the EC2 console, create a launch template that specifies the instance types that are allowed. Assign the launch template to the developers' IAM accounts.
- C. Create a new IAM policy. Specify the instance types that are allowed. Attach the policy to an IAM group that contains the IAM accounts for the developers.
- D. Use EC2 Image Builder to create an image pipeline for the developers and assist them in the creation of a golden image.

Answer: C

Explanation:

This is doable with IAM policy creation to restrict users to specific instance types. Found the below article: <https://blog.vizuri.com/limiting-allowed-aws-instance-type-with-iam-policy>

Question: 101

A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost for cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.

Which actions should a solutions architect take to resolve the problem and prevent it from

happening in the future? (Select THREE.)

- A. Create an AWS Config rule in each account to find resources with missing tags.
- B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
- C. Use Amazon Inspector in the organization to find resources with missing tags.
- D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
- E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
- F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.

Answer: A,B,E

Explanation:

<https://docs.aws.amazon.com/config/latest/developerguide/config-rule-multi-account-deployment.html>
<https://docs.aws.amazon.com/config/latest/developerguide/aggregate-data.html>
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_tagging.html

Question: 102

A company has an on-premises monitoring solution using a PostgreSQL database for persistence of events. The database is unable to scale due to heavy ingestion and it frequently runs out of storage. The company wants to create a hybrid solution and has already set up a VPN connection between its network and AWS. The solution should include the following attributes:

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

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

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

Answer: A,D

Explanation:

<https://aws.amazon.com/kinesis/data-firehose/faqs/>

Question: 103

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 an 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 are 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 blocking traffic that is responsible for high costs.
- B. Add an interface VPC endpoint for Kinesis Data Streams to the VPC. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
- 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 VPC. Ensure that the VPC endpoint policy allows traffic from the applications.

Answer: D

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html><https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/>

VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint.

Question: 104

A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1 Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.

The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that is configured to route all inter-VPC traffic within that Region. Which solution will meet these requirements?

- A. Create a private VIF from the DX-A connection into a Direct Connect gateway. Create a private VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with the Direct Connect gateway. Peer the transit gateways with each other to support cross-Region routing.
- B. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Associate the eu-west-1 transit gateway with this Direct Connect gateway. Create a transit VIF from the DX-B connection into a separate Direct Connect gateway. Associate the us-east-1 transit gateway with this separate Direct Connect gateway. Peer the Direct Connect gateways with each other to support high availability and cross-Region routing.
- C. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway. Configure the Direct Connect gateway to route traffic between the transit gateways.
- D. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway. Peer the transit gateways with each other to support cross-Region routing.

Answer: D

Explanation:

in this solution, two transit VIFs are created - one from the DX-A connection and one from the DX-B connection - into the same Direct Connect gateway for high availability. Both the eu-west-1 and us-east-1 transit gateways are then associated with this Direct Connect gateway. The transit gateways are then peered with each other to support cross-Region routing. This solution meets the requirements of the company by creating a highly available connection between the on-premises data center and the VPCs in both the eu-west-1 and us-east-1 regions, and by enabling direct traffic routing between VPCs in those regions.

Question: 105

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

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

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

Answer: A,D,E

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/send-a-notification-when-an-iam-user-is-created.html>

Question: 106

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

The company must create separate accounts for development, staging, production, and shared network. The production account and the shared network account must have connectivity to all accounts. The development account and the staging account must have access only to each other. Which combination of steps should a solutions architect take to meet these requirements? (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 Cognito user pools and identity pools to manage access to accounts and between accounts.

Answer: A,C,D

Explanation:

The correct answer would be options A, C and D, because they address the requirements outlined in the question. A. Deploying a landing zone environment using AWS Control Tower and enrolling accounts in an organization in AWS Organizations allows for a centralized management of access to all accounts and applications. C. Creating transit gateways and transit gateway VPC attachments in each account and configuring appropriate route tables allows for private network traffic, and ensures that the production account and shared network account have connectivity to all accounts, while the development and staging accounts have access only to each other. D. Setting up and enabling AWS IAM Identity Center (AWS Single Sign-On) and creating appropriate permission sets with required MFA for existing accounts allows for multi-factor authentication at login and specific roles to be assigned to user groups.

Question: 107

A company runs its application in the eu-west-1 Region and has one account for each of its environments development, testing, and production All the environments are running 24 hours a day 7 days a week by using stateful Amazon EC2 instances and Amazon RDS for MySQL databases The databases are between 500 GB and 800 GB in size

The development team and testing team work on business days during business hours, but the production environment operates 24 hours a day. 7 days a week. The company wants to reduce costs AH resources are tagged with an environment tag with either development, testing, or production as the key. What should a solutions architect do to reduce costs with the LEAST operational effort?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs once every day Configure the rule to invoke one AWS Lambda function that starts or stops instances based on the tag day and time.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening. Configure the rule to invoke an AWS Lambda function that stops instances based on the tag-Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to invoke another Lambda function that starts instances based on the tag
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening Configure the rule to invoke an AWS Lambda function that terminates instances based on the tag Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to invoke another Lambda function that restores the instances from their last backup based on the tag.
- D. Create an Amazon EventBridge rule that runs every hour. Configure the rule to invoke one AWS Lambda function that terminates or restores instances from their last backup based on the tag. day, and time.

Answer: B

Explanation:

Creating an Amazon EventBridge rule that runs every business day in the evening to stop instances and another rule that runs every business day in the morning to start instances based on the tag will reduce costs with the least operational effort. This approach allows for instances to be stopped during non-business hours when they are not in use, reducing the costs associated with running them. It also allows for instances to be started again in the morning when the development and testing teams need to use them.

Question: 108

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

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

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

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html#apig-request-throttling-account-level-limits>

Question: 109

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:

Option A, Deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC, allows the company to use its existing security tool while still running it within the AWS environment. This ensures that all packets coming in and out of the VPC are inspected by the security tool in real time. Option D, Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool, allows for high availability within an AWS Region. By provisioning a Gateway Load Balancer for each Availability Zone, the traffic is redirected to the security tool in the event of any failures or outages. This ensures that the security tool is always available to inspect the traffic, even in the event of a failure.

Question: 110

A company has purchased appliances from different vendors. The appliances all have IoT sensors. The sensors send status information in the vendors' proprietary formats to a legacy application that parses the information into JSON. The parsing is simple, but each vendor has a unique format. Once daily, the application parses all the JSON records and stores the records in a relational database for analysis.

The company needs to design a new data analysis solution that can deliver faster and optimize costs. Which solution will meet these requirements?

- A. Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the files. Use Amazon Athena and Amazon QuickSight for analysis.
- B. Migrate the application server to AWS Fargate, which will receive the information from IoT sensors and parse the information into a relational format. Save the parsed information to Amazon Redshift for analysis.
- C. Create an AWS Transfer for SFTP server. Update the IoT sensor code to send the information as a .csv file through SFTP to the server. Use AWS Glue to catalog the files. Use Amazon Athena for analysis.
- D. Use AWS Snowball Edge to collect data from the IoT sensors directly to perform local analysis. Periodically collect the data into Amazon Redshift to perform global analysis.

Answer: A

Explanation:

Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the files. Use Amazon Athena and Amazon QuickSight for analysis. This solution meets the requirement of faster analysis and cost optimization by using AWS IoT Core to collect data from the IoT sensors in real-time and then using AWS Glue and Amazon Athena for efficient data analysis.

This solution involves connecting the IoT sensors to the AWS IoT Core, setting a rule to invoke an AWS Lambda function to parse the information, and saving a .csv file to Amazon S3. AWS Glue can be used to catalog the files and Amazon Athena and Amazon QuickSight can be used for analysis. This solution will enable faster and more cost-effective data analysis.

This solution is in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that: "AWS IoT Core can be used to ingest and process the data, AWS Lambda can be used to process and transform the data, and Amazon S3 can be used to store the data. AWS Glue can be used to catalog and access the data, Amazon Athena can be used to query the data, and Amazon QuickSight can be used to visualize the data." (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf)

Question: 111

A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connection connection in a central network account.

The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.

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

- A. Create a Direct Connect gateway in the central account. In each of the accounts, create an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway.
- B. Create a Direct Connect gateway and a transit gateway in the central network account. Attach the transit gateway to the Direct Connect gateway by using a transit VIF.
- C. Provision an internet gateway. Attach the internet gateway to subnets. Allow internet traffic through the gateway.
- D. Share the transit gateway with other accounts. Attach VPCs to the transit gateway.
- E. Provision VPC peering as necessary.
- F. Provision only private subnets. Open the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center.

Answer: B,D,F

Explanation:

Option A is incorrect because creating a Direct Connect gateway in the central account and creating an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway does not enable active-passive failover between the regions. A Direct Connect gateway is a globally available resource that enables you to connect your AWS Direct Connect connection over a private virtual interface (VIF) to one or more VPCs in any AWS Region. A virtual private gateway is the VPN concentrator on the Amazon side of a VPN connection. You can associate a Direct Connect gateway with either a transit gateway or a virtual private gateway. However, a Direct Connect gateway does not provide any load balancing or failover capabilities by itself¹ Option B is correct because creating a Direct Connect gateway and a transit gateway in the central network account and attaching the transit gateway to the Direct Connect gateway by using a transit VIF meets the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. A transit VIF is a type of private VIF that you can use to connect your AWS Direct Connect connection to a transit gateway or a Direct Connect gateway. A transit gateway is a network transit hub that you can use to interconnect your VPCs and on-premises networks. By using a transit VIF, you can route traffic between your on-premises network and multiple VPCs across different AWS accounts and Regions through a single connection²³

Option C is incorrect because provisioning an internet gateway, attaching the internet gateway to subnets, and allowing internet traffic through the gateway does not meet the requirement of routing cloud resources to the internet through its on-premises data center. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. An internet gateway serves two purposes: to provide a target in your VPC route tables for internet-routable traffic, and to perform network address translation (NAT) for instances that have been assigned public IPv4 addresses. By using an internet gateway, you are routing cloud resources directly to the internet, not through your on-premises data center. Option D is correct because sharing the transit gateway with other accounts and attaching VPCs to the transit gateway meets the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. You can share your transit gateway with other AWS accounts within the same organization by using AWS Resource Access Manager (AWS RAM).

This allows you to centrally manage connectivity from multiple accounts without having to create individual peering connections between VPCs or duplicate network appliances in each account. You can attach VPCs from different accounts and Regions to your shared transit gateway and enable routing between them. Option E is incorrect because provisioning VPC peering as necessary does not meet the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. VPC peering is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single Region. However, VPC peering does not allow you to route traffic from your on-premises network to your VPCs or between multiple Regions. You would need to create multiple VPN connections or Direct Connect connections for each VPC peering connection, which increases operational complexity and costs.

Option F is correct because provisioning only private subnets, opening the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center meets the requirement of routing cloud resources to the internet through its on-premises data center. A private subnet is a subnet that's associated with a route table that has no route to an internet gateway. Instances in a private subnet can communicate with other instances in the same VPC but cannot access resources on the internet directly. To enable outbound internet access from instances in private subnets, you can use NAT devices such as NAT gateways or NAT instances that are deployed in public subnets. A public subnet is a subnet that's associated with a route table that has a route to an internet gateway. Alternatively, you can use your on-premises data center as a NAT device by configuring routes on your transit gateway and customer gateway that direct outbound internet traffic from your private subnets through your VPN connection or Direct Connect connection. This way, you can route cloud resources to the internet through your on-premises data center instead of using an internet gateway.

Reference: 1: <https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html> 2: <https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-transit-virtual-interfaces.html> 3: <https://docs.aws.amazon.com/vpc/latest/tgw/what-is-transit-gateway.html> : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Internet_Gateway.html : <https://docs.aws.amazon.com/vpc/latest/tgw/tgw-sharing.html> : <https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html> : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario2.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario3.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Instance.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Gateway.html

Question: 112

A company has hundreds of AWS accounts. The company recently implemented a centralized internal process for purchasing new Reserved Instances and modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement. Previously, business units directly purchased or modified Reserved Instances in their own respective AWS accounts autonomously.

A solutions architect needs to enforce the new process in the most secure way possible.

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

- A. Ensure that all AWS accounts are part of an organization in AWS Organizations with all features enabled.
- B. Use AWS Config to report on the attachment of an IAM policy that denies access to the `ec2:PurchaseReservedInstancesOffering` action and the `ec2:ModifyReservedInstances` action.
- C. In each AWS account, create an IAM policy that denies the

ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.

D. Create an SCP that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action. Attach the SCP to each OU of the organization.

E. Ensure that all AWS accounts are part of an organization in AWS Organizations that uses the consolidated billing feature.

Answer: A,D

Explanation:

All features - The default feature set that is available to AWS Organizations. It includes all the functionality of consolidated billing, plus advanced features that give you more control over accounts in your organization. For example, when all features are enabled the management account of the organization has full control over what member accounts can do. The management account can apply SCPs to restrict the services and actions that users (including the root user) and roles in an account can access.[https://docs.aws.amazon.com/organizations/latest/userguide/orgs_getting-](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_getting-started_concepts.html#feature-set)

[started_concepts.html#feature-set](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_getting-started_concepts.html#feature-set)

Question: 113

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:

Migrate the database to Amazon Aurora MySQL. - Create an Amazon Aurora Replica. - Use RDS Proxy in front of the database. - These options are correct because they address the requirement of reducing the failover time to less than 20 seconds. Migrating to Amazon Aurora MySQL and creating an Aurora replica can reduce the failover time to less than 20 seconds. Aurora has a built-in, fault-tolerant storage system that can automatically detect and repair failures. Additionally, Aurora has a feature called "Aurora Global Database" which allows you to create read-only replicas across multiple AWS regions which can further help to reduce the failover time. Creating an Aurora replica can also help to reduce the failover time as it can take over as the primary DB instance in case of a failure. Using RDS proxy can also help to reduce the failover time as it

can route the queries to the healthy DB instance, it also helps to balance the load across multiple DB instances.

Question: 114

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/IAM/latest/UserGuide/confused-deputy.html>
This is the most secure way to allow org1 to access resources in org2 because it allows for least privilege security access. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN) and include the external ID in the IAM role's trust policy when requesting access to perform the required tasks. This ensures that the partner company can only access the resources that it needs and only from the specific customer account.

Question: 115

A delivery company needs to migrate its third-party route planning application to AWS. The third party supplies a supported Docker image from a public registry. The image can run in as many containers as required to generate the route map.

The company has divided the delivery area into sections with supply hubs so that delivery drivers travel the shortest distance possible from the hubs to the customers. To reduce the time necessary to generate route maps, each section uses its own set of Docker containers with a custom configuration that processes orders only in the section's area.

The company needs the ability to allocate resources cost-effectively based on the number of running

containers.

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

- A. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on Amazon EC2. Use the Amazon EKS CLI to launch the planning application in pods by using the `-tags` option to assign a custom tag to the pod.
- B. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on AWS Fargate. Use the Amazon EKS CLI to launch the planning application. Use the AWS CLI `tag-resource` API call to assign a custom tag to the pod.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster on Amazon EC2. Use the AWS CLI with `run-tasks` set to `true` to launch the planning application by using the `-tags` option to assign a custom tag to the task.
- D. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate. Use the AWS CLI `run-task` command and set `enableECSTags` to `true` to launch the planning application. Use the `--tags` option to assign a custom tag to the task.

Answer: D

Explanation:

Amazon Elastic Container Service (ECS) on AWS Fargate is a fully managed service that allows you to run containers without having to manage the underlying infrastructure. When you launch tasks on Fargate, resources are automatically allocated based on the number of tasks running, which reduces the operational overhead.

Using ECS on Fargate allows you to assign custom tags to tasks using the `--tags` option in the `run-task` command, as described in the documentation: <https://docs.aws.amazon.com/cli/latest/reference/ecs/run-task.html> You can also set `enableECSTags` to `true`, which allows the service to automatically add the cluster name and service name as tags. <https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-placement-constraints.html#tag-based-scheduling>

Question: 116

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

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

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

Answer: A,E

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/11/announcing-support-for-inter-region-vpc-peering/>

Question: 117

A company's solutions architect is reviewing a new internally developed application in a sandbox AWS account. The application uses an AWS Auto Scaling group of Amazon EC2 instances that have an IAM instance profile attached. Part of the application logic creates and accesses secrets from AWS Secrets Manager. The company has an AWS Lambda function that calls the application API to test the functionality. The company also has created an AWS CloudTrail trail in the account.

The application's developer has attached the `SecretsManagerReadOnlyAccess` AWS managed IAM policy to an IAM role. The IAM role is associated with the instance profile that is attached to the EC2 instances.

The solutions architect has invoked the Lambda function for testing.

The solutions architect must replace the `SecretsManagerReadOnlyAccess` policy with a new policy that provides least privilege access to the Secrets Manager actions that the application requires. What is the MOST operationally efficient solution that meets these requirements?

- A. Generate a policy based on CloudTrail events for the IAM role. Use the generated policy output to create a new IAM policy. Use the newly generated IAM policy to replace the `SecretsManagerReadOnlyAccess` policy that is attached to the IAM role.
- B. Create an analyzer in AWS Identity and Access Management Access Analyzer. Use the IAM role's Access Advisor findings to create a new IAM policy. Use the newly created IAM policy to replace the `SecretsManagerReadOnlyAccess` policy that is attached to the IAM role.
- C. Use the `aws cloudtrail lookup-events` AWS CLI command to filter and export CloudTrail events that are related to Secrets Manager. Use a new IAM policy that contains the actions from CloudTrail to replace the `SecretsManagerReadOnlyAccess` policy that is attached to the IAM role.
- D. Use the IAM policy simulator to generate an IAM policy for the IAM role. Use the newly generated IAM policy to replace the `SecretsManagerReadOnlyAccess` policy that is attached to the IAM role.

Answer: B

Explanation:

The IAM policy simulator will generate a policy that contains only the necessary permissions for the application to access Secrets Manager, providing the least privilege necessary to get the job done. This is the most efficient solution as it will not require additional steps such as analyzing CloudTrail events or manually creating and testing an IAM policy.

You can use the IAM policy simulator to generate an IAM policy for an IAM role by specifying the role and the API actions and resources that the application or service requires. The simulator will then generate an IAM policy that grants the least privilege access to those actions and resources.

Once you have generated an IAM policy using the simulator, you can replace the existing `SecretsManagerReadOnlyAccess` policy that is attached to the IAM role with the newly generated policy. This will

ensure that the application or service has the least privilege access to the Secrets Manager actions that it requires.

You can access the IAM policy simulator through the IAM console, AWS CLI, and AWS SDKs. Here is the link for more information:https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_simulator.html

Question: 118

A solutions architect must analyze a company's Amazon EC2 Instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently. The company is running several large, high-memory EC2 instances to host database clusters that are deployed in active/passive configurations. The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern. The solutions architect must analyze the environment and take action based on the findings. Which solution meets these requirements MOST cost-effectively?

- A. Create a dashboard by using AWS Systems Manager OpsCenter. Configure visualizations for Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes. Review the dashboard periodically and identify usage patterns. Right size the EC2 instances based on the peaks in the metrics.
- B. Turn on Amazon CloudWatch detailed monitoring for the EC2 instances and their EBS volumes. Create and review a dashboard that is based on the metrics. Identify usage patterns. Right size the EC2 instances based on the peaks in the metrics.
- C. Install the Amazon CloudWatch agent on each of the EC2 instances. Turn on AWS Compute Optimizer, and let it run for at least 12 hours. Review the recommendations from Compute Optimizer, and right size the EC2 instances as directed.
- D. Sign up for the AWS Enterprise Support plan. Turn on AWS Trusted Advisor. Wait 12 hours. Review the recommendations from Trusted Advisor, and rightsize the EC2 instances as directed.

Answer: C

Explanation:

(<https://aws.amazon.com/compute-optimizer/pricing/>,<https://aws.amazon.com/systems-manager/pricing/>).

<https://aws.amazon.com/compute-optimizer/>

Question: 119

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 that is named DBA-Secret. Grant the role the required permissions to access the secrets. In the DBA account, create an IAM role that is named DBA-Admin. Grant the DBA-Admin role the required permissions to assume the DBA-Secret role in the application account. Attach the DBA-Admin role to the EC2 instance for access to the crossaccount secrets.
- 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:

Option B is correct because creating an IAM role in the application account that has permissions to access the secrets and creating an IAM role in the DBA account that has permissions to assume the role in the application account eliminates the need to manually share the secrets. This approach uses cross-account IAM roles to grant access to the secrets in the application account. The database administrators can assume the role in the application account from their EC2 instance in the DBA account and retrieve the secrets without having to store them locally or share them manually² Reference: 1: <https://docs.aws.amazon.com/ram/latest/userguide/what-is.html> 2: https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html 3: <https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html> : https://docs.aws.amazon.com/secretsmanager/latest/userguide/tutorials_basic.html : <https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html>

Question: 120

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? (Select TWO)

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

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requested-region.html

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_ec2.html

Question: 121

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.

Answer: A,C

Explanation:

<https://docs.aws.amazon.com/sns/latest/dg/sns-cross-region-delivery.html>

Question: 122

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 singlethreaded, 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 EC2 Spot Instances to run the application. Use AWS CodeDeploy to deploy and run the application every 4 hours.

Answer: C

Explanation:

step function could run a scheduled task when triggered by eventbridge, but why would you add that layer of complexity just to run aws batch when you could directly invoke it through eventbridge. The link provided - <https://aws.amazon.com/pt/blogs/compute/orchestrating-high-performance-computing-with-aws-step-functions-and-aws-batch/> makes sense only for HPC, this is a single instance that needs to be run

Question: 123

A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:

- Amazon S3 bucket that stores game assets
- Amazon DynamoDB table that stores player scores

A solutions architect needs to design a multi-Region solution that will reduce latency, improve reliability, and require the least effort to implement.

What should the solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket. Configure S3 Cross-Region Replication. Create a new DynamoDB table in a new Region. Use the new table as a replica target for DynamoDB global tables.
- B. Create an Amazon CloudFront distribution to serve assets from the S3 bucket. Configure S3 Same-Region Replication. Create a new DynamoDB table in a new Region. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC).
- C. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets. Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3

buckets in each Region. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.

D. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets- Create an Amazon CloudFront distribution and configure origin failover with two origin accessing the S3 buckets Create a new DynamoDB table in a new Region Use the new table as a replica target for DynamoDB global tables.

Answer: C

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-global-table-stream-lambda/?nc1=h_ls

Question: 124

A company has an on-premises website application that provides real estate information for potential renters and buyers. The website uses a Java backend and a NOSQL MongoDB database to store subscriber data.

The company needs to migrate the entire application to AWS with a similar structure. The application must be deployed for high availability, and the company cannot make changes to the application

Which solution will meet these requirements?

A. Use an Amazon Aurora DB cluster as the database for the subscriber data. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

B. Use MongoDB on Amazon EC2 instances as the database for the subscriber data. Deploy EC2 instances in an Auto Scaling group in a single Availability Zone for the Java backend application.

C. Configure Amazon DocumentDB (with MongoDB compatibility) with appropriately sized instances in multiple Availability Zones as the database for the subscriber data. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

D. Configure Amazon DocumentDB (with MongoDB compatibility) in on-demand capacity mode in multiple Availability Zones as the database for the subscriber data. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

Answer: C

Explanation:

On-demand capacity mode is the function of

Dynamodb.<https://aws.amazon.com/blogs/news/running-spikey-workloads-and-optimizing-costs-by-more-than-90-using-amazon-dynamodb-on-demand-capacity-mode/>

Amazon DocumentDB Elastic Clusters<https://aws.amazon.com/blogs/news/announcing-amazon-documentdb-elastic-clusters/>

Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application. This will provide high availability and scalability, while allowing the company to retain the same

database structure as the original application.

Question: 125

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:

<https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-denied-error-s3/>

Question: 126

A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data. Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days

The company has a high-speed AWS Direct Connect connection Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day

Which solution meets these requirements?

- A. Use regularly scheduled AWS Snowball Edge devices to transfer the sequencing data into AWS When AWS receives the Snowball Edge device and the data is loaded into Amazon S3 use S3 events to trigger an AWS Lambda function to process the data
- B. Use AWS Data Pipeline to transfer the sequencing data to Amazon S3 Use S3 events to trigger an Amazon EC2 Auto Scaling group to launch custom-AMI EC2 instances running the Docker containers to process the data
- C. Use AWS DataSync to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Lambda function that starts an AWS Step Functions workflow Store the Docker images in Amazon Elastic Container Registry (Amazon ECR) and trigger AWS Batch to run the container and process the sequencing data
- D. Use an AWS Storage Gateway file gateway to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Batch job that runs on Amazon EC2 instances running the Docker containers to process the data

Answer: C

Explanation:

AWS DataSync can be used to transfer the sequencing data to Amazon S3, which is a more efficient and faster method than using Snowball Edge devices. Once the data is in S3, S3 events can trigger an AWS Lambda function that starts an AWS Step Functions workflow. The Docker images can be stored in Amazon Elastic Container Registry (Amazon ECR) and AWS Batch can be used to run the container and process the sequencing data.

Question: 127

A company runs a content management application on a single Windows Amazon EC2 instance in a development environment. The application reads and writes static content to a 2 TB Amazon Elastic Block Store (Amazon EBS) volume that is attached to the instance as the root device. The company plans to deploy this application in production as a highly available and fault-tolerant solution that runs on at least three EC2 instances across multiple Availability Zones.

A solutions architect must design a solution that joins all the instances that run the application to an Active Directory domain. The solution also must implement Windows ACLs to control access to file contents. The application always must maintain exactly the same content on all running instances at any given point in time.

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

- A. Create an Amazon Elastic File System (Amazon EFS) file share. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Implement a user data script to install the application, join the instance to the AD domain, and mount the EFS file share.
- B. Create a new AMI from the current EC2 instance that is running. Create an Amazon FSx for Lustre file

system. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Implement a user data script to join the instance to the AD domain and mount the FSx for Lustre file system.

C. Create an Amazon FSx for Windows File Server file system. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Implement a user data script to install the application and mount the FSx for Windows File Server file system. Perform a seamless domain join to join the instance to the AD domain.

D. Create a new AMI from the current EC2 instance that is running. Create an Amazon Elastic File System (Amazon EFS) file system. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Perform a seamless domain join to join the instance to the AD domain.

Answer: C

Explanation:

<https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_join_instance.html

Question: 128

A software as a service (SaaS) based company provides a case management solution to customers. A part of the solution. The company uses a standalone Simple Mail Transfer Protocol (SMTP) server to send email messages from an application. The application also stores an email template for acknowledgement email messages that populate customer data before the application sends the email message to the customer.

The company plans to migrate this messaging functionality to the AWS Cloud and needs to minimize operational overhead.

Which solution will meet these requirements MOST cost-effectively?

A. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace. Store the email template in an Amazon S3 bucket. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template. Use an SDK in the Lambda function to send the email message.

B. Set up Amazon Simple Email Service (Amazon SES) to send email messages. Store the email template in an Amazon S3 bucket. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template. Use an SDK in the Lambda function to send the email message.

C. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace. Store the email template in Amazon Simple Email Service (Amazon SES) with parameters for the customer data. Create an AWS Lambda function to call the SES template and to pass customer data to replace the parameters. Use the AWS Marketplace SMTP server to send the email message.

D. Set up Amazon Simple Email Service (Amazon SES) to send email messages. Store the email template on

Amazon SES with parameters for the customer data. Create an AWS Lambda function to call the SendTemplatedEmail API operation and to pass customer data to replace the parameters and the email destination.

Answer: D

Explanation:

In this solution, the company can use Amazon SES to send email messages, which will minimize operational overhead as SES is a fully managed service that handles sending and receiving email messages. The company can store the email template on Amazon SES with parameters for the customer data and use an AWS Lambda function to call the SendTemplatedEmail API operation, passing in the customer data to replace the parameters and the email destination. This solution eliminates the need to set up and manage an SMTP server on EC2 instances, which can be costly and time-consuming.

Question: 129

A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.

The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.

Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.

How can the company solve this problem?

- A. Turn on termination protection for the EC2 instances.
- B. Update the visibility timeout for the SOS queue to 3 hours.
- C. Configure scale-in protection for the instances during processing.
- D. Update the redrive policy and set maxReceiveCount to 0.

Answer: B

Explanation:

The best solution for this problem is to update the visibility timeout for the SQS queue to 3 hours. This is because when the visibility timeout is set to 1 hour, it means that if the EC2 instance doesn't process the message within an hour, it will be moved to the dead-letter queue. By increasing the visibility timeout to 3 hours, this should give the EC2 instance enough time to process the message before it gets moved to the dead-letter queue. Additionally, configuring scale-in protection for the EC2 instances during processing will help to ensure that the instances are not terminated while the messages are being processed.

Question: 130

A company has developed APIs that use Amazon API Gateway with Regional endpoints. The APIs call AWS Lambda functions that use API Gateway authentication mechanisms. After a design review, a solutions architect identifies a set of APIs that do not require public access.

The solutions architect must design a solution to make the set of APIs accessible only from a VPC. All APIs need to be called with an authenticated user.

Which solution will meet these requirements with the LEAST amount of effort?

- A. Create an internal Application Load Balancer (ALB). Create a target group. Select the Lambda function to call. Use the ALB DNS name to call the API from the VPC.
- B. Remove the DNS entry that is associated with the API in API Gateway. Create a hosted zone in Amazon Route 53. Create a CNAME record in the hosted zone. Update the API in API Gateway with the CNAME record. Use the CNAME record to call the API from the VPC.
- C. Update the API endpoint from Regional to private in API Gateway. Create an interface VPC endpoint in the VPC. Create a resource policy, and attach it to the API. Use the VPC endpoint to call the API from the VPC.
- D. Deploy the Lambda functions inside the VPC. Provision an EC2 instance, and install an Apache server. From the Apache server, call the Lambda functions. Use the internal CNAME record of the EC2 instance to call the API from the VPC.

Answer: C

Explanation:

This solution requires the least amount of effort as it only requires to update the API endpoint to private in API Gateway and create an interface VPC endpoint. Then create a resource policy and attach it to the API. This will make the API only accessible from the VPC and still keep the authentication mechanism intact. Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/private-api-gateway-vpc-endpoint/> <https://aws.amazon.com/api-gateway/features/>

Question: 131

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:

<https://aws.amazon.com/about-aws/whats-new/2016/04/transfer-files-into-amazon-s3-up-to-300-percent-faster/>

Question: 132

A solutions architect is investigating an issue in which a company cannot establish new sessions in Amazon Workspaces. An initial analysis indicates that the issue involves user profiles. The AmazonWorkspaces environment is configured to use Amazon FSx for Windows File Server as the profile share storage. The FSx for Windows File Server file system is configured with 10 TB of storage. The solutions architect discovers that the file system has reached its maximum capacity. The solutions architect must ensure that users can regain access. The solution also must prevent the problem from occurring again.

Which solution will meet these requirements?

A. Remove old user profiles to create space. Migrate the user profiles to an Amazon FSx for Lustre file system.

B. Increase capacity by using the update-file-system command. Implement an Amazon CloudWatch metric that monitors free space. Use Amazon EventBridge to invoke an AWS Lambda function to increase capacity as required.

C. Monitor the file system by using the FreeStorageCapacity metric in Amazon CloudWatch. Use AWS Step Functions to increase the capacity as required.

D. Remove old user profiles to create space. Create an additional FSx for Windows File Server file system. Update the user profile redirection for 50% of the users to use the new file system.

Answer: B

Explanation:

It can prevent the issue from happening again by monitoring the file system with the FreeStorageCapacity metric in Amazon CloudWatch and using Amazon EventBridge to invoke an AWS Lambda function to increase the capacity as required. This ensures that the file system always has enough free space to store user profiles and avoids reaching maximum capacity.

Question: 133 An international delivery company hosts a delivery management system on AWS. Drivers use the system to upload confirmation of delivery. Confirmation includes the recipient's signature or a photo of the package with the recipient. The driver's handheld device uploads signatures and photos through FTP to a single Amazon EC2 instance. Each handheld device saves a file in a directory based on the signed-in user, and the file name matches the delivery number. The EC2 instance then adds metadata to the file after querying a central database to pull delivery information. The file is then placed in Amazon S3 for archiving.

As the company expands, drivers report that the system is rejecting connections. The FTP server is having problems because of dropped connections and memory issues. In response to these problems, a system engineer schedules a cron task to reboot the EC2 instance every 30 minutes. The billing team reports that files are not always in the archive and that the central system is not always updated.

A solutions architect needs to design a solution that maximizes scalability to ensure that the archive always receives the files and that systems are always updated. The handheld devices cannot be modified, so the company cannot deploy a new application.

Which solution will meet these requirements?

- A. Create an AMI of the existing EC2 instance. Create an Auto Scaling group of EC2 instances behind an Application Load Balancer. Configure the Auto Scaling group to have a minimum of three instances.
- B. Use AWS Transfer Family to create an FTP server that places the files in Amazon Elastic File System (Amazon EFS). Mount the EFS volume to the existing EC2 instance. Point the EC2 instance to the new path for file processing.
- C. Use AWS Transfer Family to create an FTP server that places the files in Amazon S3. Use an S3 event notification through Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda function. Configure the Lambda function to add the metadata and update the delivery system.
- D. Update the handheld devices to place the files directly in Amazon S3. Use an S3 eventnotification through Amazon Simple Queue Service (Amazon SQS) to invoke an AWS Lambda function. Configure the Lambda function to add the metadata and update the delivery system.

Answer: C

Explanation:

Using AWS Transfer Family to create an FTP server that places the files in Amazon S3 and using S3 event notifications through Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda function will ensure that the archive always receives the files and that the central system is always updated. This solution maximizes scalability and eliminates the need for manual intervention, such as rebooting the EC2 instance.

Question: 134

A company is running an application in the AWS Cloud. The application runs on containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost. Which solution will meet these requirements with the LEAST amount of operational overhead?

- A. Provision an Aurora Replica in a different Region.
- B. Set up AWS DataSync for continuous replication of the data to a different Region.
- C. Set up AWS Database Migration Service (AWS DMS) to perform a continuous replication of the data to a different Region.
- D. Use Amazon Data Lifecycle Manager (Amazon DLM) to schedule a snapshot every 5 minutes.

Answer: A

Explanation:

Provision an Aurora Replica in a different Region will meet the requirement of the application being able to recover to a separate AWS Region in the event of an application failure, and no data can be lost, with the least amount of operational overhead.

Question: 135

A financial services company receives a regular data feed from its credit card servicing partner. Approximately 5,000 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data. The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format. Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable. Which solutions will meet these requirements?

- A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue. Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
- B. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue. Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains messages. Have the application process each record, and transform the record into JSON format. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
- C. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing.
- D. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirements.

Define the output format as JSON. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Answer: C

Explanation:

You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object create events occur. The Lambda function will then trigger the Glue ETL job to transform the records masking the sensitive data and modifying the output format to JSON. This solution meets all requirements.

Question: 136

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

Which solution will achieve the company's goal with the LEAST operational overhead?

- A. Install the AWS Replication Agent on the source servers, including the MySQL servers. Set up replication for all servers. Launch test instances for regular drills. Cut over to the test instances to fail over the workload in the case of a failure event.
- B. Install the AWS Replication Agent on the source servers, including the MySQL servers. Initialize AWS Elastic Disaster Recovery in the target AWS Region. Define the launch settings. Frequently perform failover and fallback from the most recent point in time.
- C. Create AWS Database Migration Service (AWS DMS) replication servers and a target Amazon Aurora MySQL DB cluster to host the database. Create a DMS replication task to copy the existing data to the target DB cluster. Create a local AWS Schema Conversion Tool (AWS SCT) change data capture (CDC) task to keep the data synchronized. Install the rest of the software on EC2 instances by starting with a compatible base AMI.
- D. Deploy an AWS Storage Gateway Volume Gateway on premises. Mount volumes on all on-premises servers. Install the application and the MySQL database on the new volumes. Take regular snapshots. Install all the software on EC2 instances by starting with a compatible base AMI. Launch a Volume Gateway on an EC2 instance. Restore the volumes from the latest snapshot. Mount the new volumes on the EC2 instances in the case of a failure event.

Answer: B

Explanation:

<https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html>
<https://docs.aws.amazon.com/drs/latest/userguide/recovery-workflow-gs.html>

Question: 137

A company is subject to regulatory audits of its financial information. External auditors who use a single AWS

account need access to the company's AWS account. A solutions architect must provide the auditors with secure, read-only access to the company's AWS account. The solution must comply with AWS security best practices.

Which solution will meet these requirements?

- A. In the company's AWS account, create resource policies for all resources in the account to grant access to the auditors' AWS account. Assign a unique external ID to the resource policy.
- B. In the company's AWS account create an IAM role that trusts the auditors' AWS account Create an IAM policy that has the required permissions. Attach the policy to the role. Assign a unique external ID to the role's trust policy.
- C. In the company's AWS account, create an IAM user. Attach the required IAM policies to the IAM user. Create API access keys for the IAM user. Share the access keys with the auditors.
- D. In the company's AWS account, create an IAM group that has the required permissions Create an IAM user in the company's account for each auditor. Add the IAM users to the IAM group.

Answer: B

Explanation:

This solution will allow the external auditors to have read-only access to the company's AWS account while being compliant with AWS security best practices. By creating an IAM role, which is a secure and flexible way of granting access to AWS resources, and trusting the auditors' AWS account, the company can ensure that the auditors only have the permissions that are required for their role and nothing more. Assigning a unique external ID to the role's trust policy, it will ensure that only the auditors' AWS account can assume the role.

Reference:

AWS IAM Roles documentation:<https://aws.amazon.com/iam/features/roles/>

AWS IAM Best practices:<https://aws.amazon.com/iam/security-best-practices/>

Question: 138

A company has a latency-sensitive trading platform that uses Amazon DynamoDB as a storage backend. The company configured the DynamoDB table to use on-demand capacity mode. A solutions architect needs to design a solution to improve the performance of the trading platform. The new solution must ensure high availability for the trading platform.

Which solution will meet these requirements with the LEAST latency?

- A. Create a two-node DynamoDB Accelerator (DAX) cluster Configure an application to read and write data by using DAX.
- B. Create a three-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.
- C. Create a three-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data directly

from the DynamoDB table and to write data by using DAX.

D. Create a single-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.

Answer: B

Explanation:

A DAX cluster can be deployed with one or two nodes for development or test workloads. One- and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data. A DAX cluster can be deployed with one or two nodes for development or test workloads. One- and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.cluster.html>

Question: 139

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

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

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

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

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

Answer: B,D

Explanation:

Moving the application frontend to a static website that is hosted on Amazon S3 will save cost as S3 is cheaper than running EC2 instances.

Using Spot instances for the backend EC2 instances will also save cost, as they are significantly cheaper than On-Demand instances. This will be suitable for the application, as it has minimal traffic during the rest of the day, and the availability of spot instances will not negatively affect the application's availability.

Reference:

Amazon S3 pricing:<https://aws.amazon.com/s3/pricing/>

Amazon EC2 Spot Instances documentation:<https://aws.amazon.com/ec2/spot/> AWS Elastic Beanstalk documentation:<https://aws.amazon.com/elasticbeanstalk/>

Amazon Elastic Compute Cloud (EC2) pricing:<https://aws.amazon.com/ec2/pricing/>

Question: 140

A company is running an event ticketing platform on AWS and wants to optimize the platform's cost-effectiveness. The platform is deployed on Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 and is backed by an Amazon RDS for MySQL DB instance. The company is developing new application features to run on Amazon EKS with AWS Fargate.

The platform experiences infrequent high peaks in demand. The surges in demand depend on event dates.

Which solution will provide the MOST cost-effective setup for the platform?

- A. Purchase Standard Reserved Instances for the EC2 instances that the EKS cluster uses in its baseline load. Scale the cluster with Spot Instances to handle peaks. Purchase 1-year All Upfront Reserved Instances for the database to meet predicted peak load for the year.
 - B. Purchase Compute Savings Plans for the predicted medium load of the EKS cluster. Scale the cluster with On-Demand Capacity Reservations based on event dates for peaks. Purchase 1-year No Upfront Reserved Instances for the database to meet the predicted base load. Temporarily scale out database read replicas during peaks.
 - C. Purchase EC2 Instance Savings Plans for the predicted base load of the EKS cluster. Scale the cluster with Spot Instances to handle peaks. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base load. Temporarily scale up the DB instance manually during peaks.
 - D. Purchase Compute Savings Plans for the predicted base load of the EKS cluster. Scale the cluster with Spot Instances to handle peaks. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base load. Temporarily scale up the DB instance manually during peaks.
-

Answer: B

Explanation:

They all mention using spot instances and EKS based on EC2. A spot instance is not appropriate for a production server and the company is developing new application designed for AWS Fargate, which means we must plan the future cost improvement including AWS Fargate.
<https://aws.amazon.com/savingsplans/compute-pricing/>

Question: 141

A company has deployed an application on AWS Elastic Beanstalk. The application uses Amazon Aurora for the database layer. An Amazon CloudFront distribution serves web requests and includes the Elastic Beanstalk domain name as the origin server. The distribution is configured with an alternate domain name that visitors use when they access the application.

Each week, the company takes the application out of service for routine maintenance. During the time that the application is unavailable, the company wants visitors to receive an informational message instead of a CloudFront error message.

A solutions architect creates an Amazon S3 bucket as the first step in the process.

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

- A. Upload static informational content to the S3 bucket.
- B. Create a new CloudFront distribution. Set the S3 bucket as the origin.
- C. Set the S3 bucket as a second origin in the original CloudFront distribution. Configure the distribution and the S3 bucket to use an origin access identity (OAI).
- D. During the weekly maintenance, edit the default cache behavior to use the S3 origin. Revert the change when the maintenance is complete.
- E. During the weekly maintenance, create a cache behavior for the S3 origin on the new distribution. Set the path pattern to \ Set the precedence to 0. Delete the cache behavior when the maintenance is complete.
- F. During the weekly maintenance, configure Elastic Beanstalk to serve traffic from the S3 bucket.

Answer: A,C,D

Explanation:

The company wants to serve static content from an S3 bucket during the maintenance period. To do this, the following steps are required:

Upload static informational content to the S3 bucket. This will provide the source of the content that will be served to the visitors.

Set the S3 bucket as a second origin in the original CloudFront distribution. Configure the distribution and the S3 bucket to use an origin access identity (OAI). This will allow CloudFront to access the S3 bucket securely and prevent public access to the bucket.

During the weekly maintenance, edit the default cache behavior to use the S3 origin. Revert the change when

the maintenance is complete. This will redirect all web requests to the S3 bucket instead of the Elastic Beanstalk domain name.

The other options are not correct because:

Creating a new CloudFront distribution is not necessary and would require changing the alternate domain name configuration.

Creating a cache behavior for the S3 origin on a new distribution would not work because the visitors would still access the original distribution using the alternate domain name.

Configuring Elastic Beanstalk to serve traffic from the S3 bucket is not possible and would not achieve the desired result.

Reference:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-values-specify.html#DownloadDistValuesPathPattern>

Question: 142

A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.

The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output. The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.

A solutions architect needs to simplify this process to minimize disruption to users.

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

- A. Directly modify the environment variables of the published Lambda function version. Use the SLATEST version to test image processing parameters.
- B. Create an Amazon DynamoDB table to store the image processing parameters. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.
- C. Directly code the image processing parameters within the Lambda function and remove the environment variables. Publish a new function version when the company updates the parameters.
- D. Create a Lambda function alias. Modify the client application to use the function alias ARN. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Answer: D

Explanation:

A Lambda function alias allows you to point to a specific version of a function and also can be updated to point

to a new version of the function without modifying the client application. This way, the company can test different versions of the function with different environment variables and, once the optimal parameters are found, update the alias to point to the new version, without the need to update the client application. By using this approach, the company can simplify the process of updating the environment variables, minimize disruption to users, and reduce the operational overhead.

Reference:

AWS Lambda documentation:<https://aws.amazon.com/lambda/>

AWS Lambda Aliases documentation:<https://docs.aws.amazon.com/lambda/latest/dg/aliases-intro.html>

AWS Lambda versioning and aliases

documentation:<https://aws.amazon.com/blogs/compute/versioning-aliases-in-aws-lambda/>

Question: 143

A global media company is planning a multi-Region deployment of an application. Amazon DynamoDB global tables will back the deployment to keep the user experience consistent across the two continents where users are concentrated. Each deployment will have a public Application Load Balancer (ALB). The company manages public DNS internally. The company wants to make the application available through an apex domain.

Which solution will meet these requirements with the LEAST effort?

- A. Migrate public DNS to Amazon Route 53. Create CNAME records for the apex domain to point to the ALB. Use a geolocation routing policy to route traffic based on user location.
- B. Place a Network Load Balancer (NLB) in front of the ALB. Migrate public DNS to Amazon Route 53. Create a CNAME record for the apex domain to point to the NLB's static IP address. Use a geolocation routing policy to route traffic based on user location.
- C. Create an AWS Global Accelerator accelerator with multiple endpoint groups that target endpoints in appropriate AWS Regions. Use the accelerator's static IP address to create a record in public DNS for the apex domain.
- D. Create an Amazon API Gateway API that is backed by AWS Lambda in one of the AWS Regions. Configure a Lambda function to route traffic to application deployments by using the round robin method. Create CNAME records for the apex domain to point to the API's URL.

Answer: C

Explanation:

AWS Global Accelerator is a service that directs traffic to optimal endpoints (in this case, the Application Load Balancer) based on the health of the endpoints and network routing. It allows you to create an accelerator that directs traffic to multiple endpoint groups, one for each Region where the application is deployed. The accelerator uses the AWS global network to optimize the traffic routing to the healthy endpoint.

By using Global Accelerator, the company can use a single static IP address for the apex domain, and traffic will be directed to the optimal endpoint based on the user's location, without the need for additional load balancers or routing policies.

Reference:

AWS Global Accelerator documentation:<https://aws.amazon.com/global-accelerator/Routing-User-Traffic>

to the Optimal AWS Region using Global Accelerator

documentation:<https://aws.amazon.com/blogs/networking-and-content-delivery/routing-user-traffic-to-the-optimal-aws-region-using-global-accelerator/>

Question: 144

A company is developing a new serverless API by using Amazon API Gateway and AWS Lambda. The company integrated the Lambda functions with API Gateway to use several shared libraries and custom classes.

A solutions architect needs to simplify the deployment of the solution and optimize for code reuse.

Which solution will meet these requirements?

- A. Deploy the shared libraries and custom classes into a Docker image. Store the image in an S3 bucket. Create a Lambda layer that uses the Docker image as the source. Deploy the API's Lambda functions as Zip packages. Configure the packages to use the Lambda layer.
- B. Deploy the shared libraries and custom classes to a Docker image. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Create a Lambda layer that uses the Docker image as the source. Deploy the API's Lambda functions as Zip packages. Configure the packages to use the Lambda layer.
- C. Deploy the shared libraries and custom classes to a Docker container in Amazon Elastic Container Service (Amazon ECS) by using the AWS Fargate launch type. Deploy the API's Lambda functions as Zip packages. Configure the packages to use the deployed container as a Lambda layer.
- D. Deploy the shared libraries, custom classes, and code for the API's Lambda functions to a Docker image. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Configure the API's Lambda functions to use the Docker image as the deployment package.

Answer: B

Explanation:

Deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (Amazon ECR) and creating a Lambda layer that uses the Docker image as the source. Then, deploying the API's Lambda functions as Zip packages and configuring the packages to use the Lambda layer would meet the requirements for simplifying the deployment and optimizing for code reuse.

A Lambda layer is a distribution mechanism for libraries, custom runtimes, and other function dependencies. It allows you to manage your in-development function code separately from your dependencies, this way you can easily update your dependencies without having to update your entire function code.

By deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (ECR), it makes it easy to manage and version the dependencies. This way, the company can use the same version of the dependencies across different Lambda functions.

By creating a Lambda layer that uses the Docker image as the source, the company can configure the API's Lambda functions to use the layer, reducing the need to include the dependencies in each function package, and making it easy to update the dependencies across all functions at once. Reference:

AWS Lambda Layers documentation:<https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html>

AWS Elastic Container Registry (ECR) documentation:<https://aws.amazon.com/ecr/Building-Lambda-Layers>

with Docker documentation:<https://aws.amazon.com/blogs/compute/building-lambda-layers-with-docker/>

Question: 145

A manufacturing company is building an inspection solution for its factory. The company has IP cameras at the end of each assembly line. The company has used Amazon SageMaker to train a machine learning (ML) model to identify common defects from still images.

The company wants to provide local feedback to factory workers when a defect is detected. The company must be able to provide this feedback even if the factory's internet connectivity is down. The company has a local Linux server that hosts an API that provides local feedback to the workers. How should the company deploy the ML model to meet these requirements?

A. Set up an Amazon Kinesis video stream from each IP camera to AWS. Use Amazon EC2 instances to take still images of the streams. Upload the images to an Amazon S3 bucket. Deploy a SageMaker endpoint with the ML model. Invoke an AWS Lambda function to call the inference endpoint when new images are uploaded.

Configure the Lambda function to call the local API when a defect is detected.

B. Deploy AWS IoT Greengrass on the local server. Deploy the ML model to the Greengrass server. Create a Greengrass component to take still images from the cameras and run inference. Configure the component to call the local API when a defect is detected.

C. Order an AWS Snowball device. Deploy a SageMaker endpoint the ML model and an Amazon EC2 instance on the Snowball device. Take still images from the cameras. Run inference from the EC2 instance. Configure the instance to call the local API when a defect is detected.

D. Deploy Amazon Monitron devices on each IP camera. Deploy an Amazon Monitron Gateway on premises. Deploy the ML model to the Amazon Monitron devices. Use Amazon Monitron health state alarms to call the local API from an AWS Lambda function when a defect is detected.

Answer: B

Explanation: The company should use AWS IoT Greengrass to deploy the ML model to the local server and provide local feedback to the factory workers. AWS IoT Greengrass is a service that extends AWS cloud capabilities to local devices, allowing them to collect and analyze data closer to the source of information, react autonomously to local events, and communicate securely with each other on local networks¹. AWS IoT Greengrass also supports ML inference at the edge, enabling devices to run ML models locally without requiring internet connectivity².

The other options are not correct because:

Setting up an Amazon Kinesis video stream from each IP camera to AWS would not work if the factory's internet connectivity is down. It would also incur unnecessary costs and latency to stream video data to the cloud and back.

Ordering an AWS Snowball device would not be a scalable or cost-effective solution for deploying the ML model. AWS Snowball is a service that provides physical devices for data transfer and edge computing, but it is not designed for continuous operation or frequent updates³.

Deploying Amazon Monitron devices on each IP camera would not work because Amazon Monitron is a service that monitors the condition and performance of industrial equipment using sensors and machine learning, not cameras⁴.

Reference: <https://aws.amazon.com/greengrass/>
<https://docs.aws.amazon.com/greengrass/v2/developerguide/use-machine-learning-inference.html><https://aws.amazon.com/snowball/> <https://aws.amazon.com/monitron/>

Question: 146

A solutions architect must create a business case for migration of a company's on-premises data center to the AWS Cloud. The solutions architect will use a configuration management database (CMDB) export of all the company's servers to create the case.

Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS Well-Architected Tool to import the CMDB data to perform an analysis and generate recommendations.
- B. Use Migration Evaluator to perform an analysis. Use the data import template to upload the data from the CMDB export.
- C. Implement resource matching rules. Use the CMDB export and the AWS Price List Bulk API to query CMDB data against AWS services in bulk.
- D. Use AWS Application Discovery Service to import the CMDB data to perform an analysis.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/architecture/accelerating-your-migration-to-aws/> Build a business case with AWS Migration Evaluator The foundation for a successful migration starts with a defined business objective (for example, growth or new offerings). In order to enable the business drivers, the established business case must then be aligned to a technical capability (increased security and elasticity). AWS Migration Evaluator (formerly known as TSO Logic) can help you meet these objectives. To get started, you can choose to upload exports from third-party tools such as Configuration Management Database (CMDB) or install a collector agent to monitor. You will receive an assessment after data collection, which includes a projected cost estimate and savings of running your on-premises workloads in the AWS Cloud. This estimate will provide a summary of the projected costs to re-host on AWS based on usage patterns. It will show the breakdown of costs by infrastructure and software licenses. With this information, you can make the business case and plan next steps.

Question: 147

A company has a website that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB is associated with an AWS WAF web ACL.

The website often encounters attacks in the application layer. The attacks produce sudden and significant increases in traffic on the application server. The access logs show that each attack originates from different IP addresses. A solutions architect needs to implement a solution to mitigate these attacks.

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

- A. Create an Amazon CloudWatch alarm that monitors server access. Set a threshold based on access by IP address. Configure an alarm action that adds the IP address to the web ACL's deny list.
- B. Deploy AWS Shield Advanced in addition to AWS WAF. Add the ALB as a protected resource.
- C. Create an Amazon CloudWatch alarm that monitors user IP addresses. Set a threshold based on access by IP address. Configure the alarm to invoke an AWS Lambda function to add a deny rule in the application server's subnet route table for any IP addresses that activate the alarm.
- D. Inspect access logs to find a pattern of IP addresses that launched the attacks. Use an Amazon Route 53 geolocation routing policy to deny traffic from the countries that host those IP addresses.

Answer: C

Explanation:

"The AWS WAF API supports security automation such as blacklisting IP addresses that exceed request limits, which can be useful for mitigating HTTP flood attacks."

><https://aws.amazon.com/blogs/security/how-to-protect-dynamic-web-applications-against-ddos-attacks-by-using-amazon-cloudfront-and-amazon-route-53/>

Question: 148

A company has a critical application in which the data tier is deployed in a single AWS Region. The data tier uses an Amazon DynamoDB table and an Amazon Aurora MySQL DB cluster. The current Aurora MySQL engine version supports a global database. The application tier is already deployed in two Regions.

Company policy states that critical applications must have application tier components and data tier components deployed across two Regions. The RTO and RPO must be no more than a few minutes each. A solutions architect must recommend a solution to make the data tier compliant with company policy.

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

- A. Add another Region to the Aurora MySQL DB cluster
- B. Add another Region to each table in the Aurora MySQL DB cluster
- C. Set up scheduled cross-Region backups for the DynamoDB table and the Aurora MySQL DB cluster
- D. Convert the existing DynamoDB table to a global table by adding another Region to its configuration
- E. Use Amazon Route 53 Application Recovery Controller to automate database backup and recovery to the secondary Region

Answer: A,D

Explanation:

The company should use Amazon Aurora global database and Amazon DynamoDB global table to deploy the data tier components across two Regions. Amazon Aurora global database is a feature that allows a single Aurora database to span multiple AWS Regions, enabling low-latency global reads and fast recovery from Region-wide outages¹. Amazon DynamoDB global table is a feature that allows a single DynamoDB table to

span multiple AWS Regions, enabling low-latency global reads and writes and fast recovery from Region-wide outages².

Reference: <https://aws.amazon.com/rds/aurora/global-database/>

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables_HowItWorks.html

<https://aws.amazon.com/route53/application-recovery-controller/>

Question: 149

A telecommunications company is running an application on AWS. The company has set up an AWS Direct Connect connection between the company's on-premises data center and AWS. The company deployed the application on Amazon EC2 instances in multiple Availability Zones behind an internal Application Load Balancer (ALB). The company's clients connect from the on-premises network by using HTTPS. The TLS terminates in the ALB. The company has multiple target groups and uses pathbased routing to forward requests based on the URL path.

The company is planning to deploy an on-premises firewall appliance with an allow list that is based on IP address. A solutions architect must develop a solution to allow traffic flow to AWS from the onpremises network so that the clients can continue to access the application.

Which solution will meet these requirements?

- A. Configure the existing ALB to use static IP addresses. Assign IP addresses in multiple Availability Zones to the ALB. Add the ALB IP addresses to the firewall appliance.
- B. Create a Network Load Balancer (NLB). Associate the NLB with one static IP addresses in multiple Availability Zones. Create an ALB-type target group for the NLB and add the existing ALB IP addresses to the firewall appliance. Update the clients to connect to the NLB.
- C. Create a Network Load Balancer (NLB). Associate the LNB with one static IP addresses in multiple Availability Zones. Add the existing target groups to the NLB. Update the clients to connect to the NLB. Delete the ALB Add the NLB IP addresses to the firewall appliance.
- D. Create a Gateway Load Balancer (GWLB). Assign static IP addresses to the GWLB in multiple Availability Zones. Create an ALB-type target group for the GWLB and add the existing ALB. Add the GWLB IP addresses to the firewall appliance. Update the clients to connect to the GWLB.

Answer: B

Explanation:

The company should create a Network Load Balancer (NLB) and associate it with one static IP address in multiple Availability Zones. The company should also create an ALB-type target group for the NLB and add the existing ALB. The company should add the NLB IP addresses to the firewall appliance and update the clients to connect to the NLB. This solution will allow traffic flow to AWS from the onpremises network by using static IP addresses that can be added to the firewall appliance's allow list. The NLB will forward requests to the ALB, which will use path-based routing to forward requests to the target groups.

Question: 150

A company runs an application on a fleet of Amazon EC2 instances that are in private subnets behind an internet-facing Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. An AWS WAF web ACL that contains various AWS managed rules is associated with the CloudFront distribution.

The company needs a solution that will prevent internet traffic from directly accessing the ALB. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new web ACL that contains the same rules that the existing web ACL contains. Associate the new web ACL with the ALB.
- B. Associate the existing web ACL with the ALB.
- C. Add a security group rule to the ALB to allow traffic from the AWS managed prefix list for CloudFront only.
- D. Add a security group rule to the ALB to allow only the various CloudFront IP address ranges.

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2022/02/amazon-cloudfront-managed-prefix-list/>

Question: 151

A company is running an application that uses an Amazon ElastiCache for Redis cluster as a caching layer. A recent security audit revealed that the company has configured encryption at rest for ElastiCache. However, the company did not configure ElastiCache to use encryption in transit. Additionally, users can access the cache without authentication.

A solutions architect must make changes to require user authentication and to ensure that the company is using end-to-end encryption. Which solution will meet these requirements?

- A. Create an AUTH token. Store the token in AWS System Manager Parameter Store, as an encrypted parameter. Create a new cluster with AUTH and configure encryption in transit. Update the application to retrieve the AUTH token from Parameter Store when necessary and to use the AUTH token for authentication.
- B. Create an AUTH token. Store the token in AWS Secrets Manager. Configure the existing cluster to use the AUTH token and configure encryption in transit. Update the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication.
- C. Create an SSL certificate. Store the certificate in AWS Secrets Manager. Create a new cluster and configure encryption in transit. Update the application to retrieve the SSL certificate from Secrets Manager when necessary and to use the certificate for authentication.
- D. Create an SSL certificate. Store the certificate in AWS Systems Manager Parameter Store, as an encrypted advanced parameter. Update the existing cluster to configure encryption in transit. Update the application to retrieve the SSL certificate from Parameter Store when necessary and to use the certificate for authentication.

Answer: B

Explanation:

Creating an AUTH token and storing it in AWS Secrets Manager and configuring the existing cluster to use the AUTH token and configure encryption in transit, and updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, would meet the requirements for user authentication and end-to-end encryption.

AWS Secrets Manager is a service that enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Secrets Manager also enables you to encrypt the data and ensure that only authorized users and applications can access it. By configuring the existing cluster to use the AUTH token and encryption in transit, all data will be encrypted as it is sent over the network, providing additional security for the data stored in ElastiCache.

Additionally, by updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, it ensures that only authorized users and applications can access the cache.

Reference:

AWS Secrets Manager documentation:<https://aws.amazon.com/secrets-manager/>

Encryption in transit for ElastiCache:<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/encryption.html>

Authentication and Authorization for

ElastiCache:<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/accessing-elasticache.html>

Question: 152

A company is running a compute workload by using Amazon EC2 Spot Instances that are in an Auto Scaling group. The launch template uses two placement groups and a single instance type.

Recently, a monitoring system reported Auto Scaling instance launch failures that correlated with longer wait times for system users. The company needs to improve the overall reliability of the workload.

Which solution will meet this requirement?

- A. Replace the launch template with a launch configuration to use an Auto Scaling group that uses attribute-based instance type selection.
- B. Create a new launch template version that uses attribute-based instance type selection. Configure the Auto Scaling group to use the new launch template version.
- C. Update the launch template Auto Scaling group to increase the number of placement groups.
- D. Update the launch template to use a larger instance type.

Answer: B

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/create-asg-instance-type->

requirements.html#use-attribute-based-instance-type-selection-prerequisites

Question: 153

A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.

During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.

Which solution will meet these requirements with the LEAST amount of effort?

- A. Migrate the application to an AWS Lambda function. Use the AWS SDK for Java to generate, modify, and access the files that the company stores directly in Amazon S3.
- B. Set up an Amazon S3 File Gateway and configure a file share that is linked to the document store. Mount the file share on an Amazon EC2 instance by using NFS. When changes occur in Amazon S3, initiate a RefreshCache API call to update the S3 File Gateway.
- C. Configure Amazon FSx for Lustre with an import and export policy. Link the new file system to an S3 bucket. Install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS.
- D. Configure AWS DataSync to connect to an Amazon EC2 instance. Configure a task to synchronize the generated files to and from Amazon S3.

Answer: C

Explanation:

The company should configure Amazon FSx for Lustre with an import and export policy. The company should link the new file system to an S3 bucket. The company should install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS. This solution will meet the requirements with the least amount of effort because Amazon FSx for Lustre is a fully managed service that provides a high-performance filesystem optimized for fast processing of workloads such as machine learning, highperformance computing, video processing, financial modeling, and electronic design automation¹. Amazon FSx for Lustre can be linked to an S3 bucket and can import data from and export data to the bucket². The import and export policy can be configured to automatically import new or changed objects from S3 and export new or changed files to S3³. This will ensure that the files are available to the public for download within 30 minutes. Amazon FSx for Lustre supports NFS version 3.0 protocol for Linux clients.

The other options are not correct because:

Migrating the application to an AWS Lambda function would require a lot of effort and may not be feasible for the existing server that generates many documents. Lambda functions have limitations on execution time, memory, disk space, and network bandwidth.

Setting up an Amazon S3 File Gateway would not work because S3 File Gateway does not support write-back caching, which means that files written to the file share are uploaded to S3 immediately and are not available locally until they are downloaded again. This would not provide fast local access to the files that the server generates and modifies.

Configuring AWS DataSync to connect to an Amazon EC2 instance would not meet the requirement of making the files available to the public for download within 30 minutes. DataSync is a service that transfers data between on-premises storage systems and AWS storage services over the internet or AWS Direct Connect. DataSync tasks can be scheduled to run at specific times or intervals, but they are not triggered by file changes.

Reference: <https://aws.amazon.com/fsx/lustre/>

<https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html>

<https://docs.aws.amazon.com/fsx/latest/LustreGuide/import-export-data-repositories.html>

<https://docs.aws.amazon.com/fsx/latest/LustreGuide/mounting-on-premises.html>

<https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html>

<https://docs.aws.amazon.com/storagegateway/latest/userguide/StorageGatewayConcepts.html>

<https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html>

Question: 154

A company is running a web application in a VPC. The web application runs on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is using AWS WAF.

An external customer needs to connect to the web application. The company must provide IP addresses to all external customers.

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

- A. Replace the ALB with a Network Load Balancer (NLB). Assign an Elastic IP address to the NLB.
- B. Allocate an Elastic IP address. Assign the Elastic IP address to the ALB. Provide the Elastic IP address to the customer.
- C. Create an AWS Global Accelerator standard accelerator. Specify the ALB as the accelerator's endpoint. Provide the accelerator's IP addresses to the customer.
- D. Configure an Amazon CloudFront distribution. Set the ALB as the origin. Ping the distribution's DNS name to determine the distribution's public IP address. Provide the IP address to the customer.

Answer: C

Explanation:

<https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html>

Option A is wrong. AWS WAF does not support associating with NLB.

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html> Option B is wrong. An ALB does not support an Elastic IP address. <https://aws.amazon.com/elasticloadbalancing/features/>

Question: 155

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.)

- www.atmicnetworks.com www.atmicnetworks.com www.atmicnetworks.com www.atmicnetworks.com
- A. Use AWS CloudFormation StackSets to deploy AWS Config rules on production accounts.
 - B. Create a new AWS Control Tower landing zone in an existing developer account. Create OUs for accounts. Add production and development accounts to production and development OUs, **respectively**.
 - C. Create a new AWS Control Tower landing zone in the company's management account. Add production and development accounts to production and development OUs, **respectively**.
 - D. Invite existing accounts to join the organization in AWS Organizations. Create SCPs to ensure **compliance**.
 - E. Create a guardrail from the management account to detect EBS encryption.
 - F. Create a guardrail for the production OU to detect EBS encryption.

Answer: C,D,F

Explanation:

<https://docs.aws.amazon.com/controltower/latest/userguide/controls.html>
<https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#ebs-enable-encryption> AWS is now transitioning the previous term 'guardrail' new term 'control'.

Question: 156 A company is running a critical stateful web application on two Linux Amazon EC2 instances behind an Application Load Balancer (ALB) with an Amazon RDS for MySQL database. The company hosts the DNS records for the application in Amazon Route 53. A solutions architect must recommend a **solution to improve the resiliency of the application**.

The solution must meet the following objectives:

- Application tier RPO of 2 minutes. RTO of 30 minutes
- Database tier RPO of 5 minutes. RTO of 30 minutes

The company does not want to make significant changes to the existing application architecture. The company **must ensure optimal latency after a failover**.

Which solution will meet these requirements?

- A. Configure the EC2 instances to use AWS Elastic Disaster Recovery. Create a cross-Region read replica for the RDS DB instance. Create an ALB in a second AWS Region. Create an AWS Global Accelerator endpoint and associate the endpoint with the ALBs. Update DNS records to point to the **Global Accelerator endpoint**.
- B. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes. Configure RDS automated backups. Configure backup replication to a second AWS Region. Create an ALB in the second Region. Create an AWS Global Accelerator endpoint, and associate the endpoint with the ALBs. Update DNS records to point to the **Global Accelerator endpoint**.
- C. Create a backup plan in AWS Backup for the EC2 instances and RDS DB instance. Configure backup replication to a second AWS Region. Create an ALB in the second Region. Configure an Amazon CloudFront distribution in front of the ALB. Update DNS records to point to CloudFront.
- D. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes. Create a cross-Region read replica for the RDS DB instance. Create an ALB in a second AWS Region. Create an AWS Global Accelerator endpoint and associate the endpoint **with the ALBs**.

Answer: B

Explanation:

This option meets the RPO and RTO requirements for both the application and database tiers and uses tools like Amazon DLM and RDS automated backups to create and manage the backups. Additionally, it uses Global Accelerator to ensure low latency after failover by directing traffic to the closest healthy endpoint.

Question: 157

A company uses an AWS CodeCommit repository. The company must store a backup copy of the data that is in the repository in a second AWS Region.

Which solution will meet these requirements?

- A. Configure AWS Elastic Disaster Recovery to replicate the CodeCommit repository data to the second Region.
- B. Use AWS Backup to back up the CodeCommit repository on an hourly schedule. Create a cross-Region copy in the second Region.
- C. Create an Amazon EventBridge rule to invoke AWS CodeBuild when the company pushes code to the repository. Use CodeBuild to clone the repository. Create a zip file of the content. Copy the file to an S3 bucket in the second Region.
- D. Create an AWS Step Functions workflow on an hourly schedule to take a snapshot of the CodeCommit repository. Configure the workflow to copy the snapshot to an S3 bucket in the second Region.

Answer: B

Explanation:

AWS Backup is a fully managed service that makes it easy to centralize and automate the creation, retention, and restoration of backups across AWS services. It provides a way to schedule automatic backups for CodeCommit repositories on an hourly basis. Additionally, it also supports cross-Region replication, which allows you to copy the backups to a second Region for disaster recovery.

By using AWS Backup, the company can set up an automatic and regular backup schedule for the CodeCommit repository, ensuring that the data is regularly backed up and stored in a second Region. This can provide a way to recover quickly from any disaster event that might occur.

Reference:

AWS Backup documentation: <https://aws.amazon.com/backup/>

AWS Backup for AWS CodeCommit documentation: <https://aws.amazon.com/about-aws/whats-new/2020/07/aws-backup-now-supports-aws-codecommit-repositories/>

Question: 158

A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company's marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.

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

- A. Instruct each business unit to add a unique secondary CIDR range to the business unit's VPC. Peer the VPCs and use a private NAT gateway in the secondary range to route traffic to the marketing team.
- B. Create an Amazon EC2 instance to serve as a virtual appliance in the marketing account's VPC. Create an AWS Site-to-Site VPN connection between the marketing team and each business unit's VPC. Perform NAT where necessary.
- C. Create an AWS PrivateLink endpoint service to share the marketing application. Grant permission to specific AWS accounts to connect to the service. Create interface VPC endpoints in other accounts to access the application by using private IP addresses.
- D. Create a Network Load Balancer (NLB) in front of the marketing application in a private subnet. Create an API Gateway API. Use the Amazon API Gateway private integration to connect the API to the NLB. Activate IAM authorization for the API. Grant access to the accounts of the other business units.

Answer: C

Explanation:

With AWS PrivateLink, the marketing team can create an endpoint service to share their internal application with other accounts securely using private IP addresses. They can grant permission to specific AWS accounts to connect to the service and create interface VPC endpoints in the other accounts to access the application by using private IP addresses. This option does not require any changes to the network of the other business units, and it does not require peering or NATing. This solution is both scalable and secure.

<https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/>

Question: 159

A company needs to audit the security posture of a newly acquired AWS account. The company's data security team requires a notification only when an Amazon S3 bucket becomes publicly exposed. The company has already established an Amazon Simple Notification Service (Amazon SNS) topic that has the data security team's email address subscribed.

Which solution will meet these requirements?

- A. Create an S3 event notification on all S3 buckets for the isPublic event. Select the SNS topic as the target for the event notifications.
- B. Create an analyzer in AWS Identity and Access Management Access Analyzer. Create an Amazon EventBridge rule for the event type "Access Analyzer Finding" with a filter for "isPublic: true." Select the SNS topic as the EventBridge rule target.
- C. Create an Amazon EventBridge rule for the event type "Bucket-Level API Call via CloudTrail" with a filter for "PutBucketPolicy." Select the SNS topic as the EventBridge rule target.
- D. Activate AWS Config and add the cloudtrail-s3-dataevents-enabled rule. Create an Amazon EventBridge rule for the event type "Config Rules Re-evaluation Status" with a filter for "NON_COMPLIANT." Select the SNS topic as the EventBridge rule target.

Answer: B

Explanation:

Access Analyzer is to assess the access policy. https://docs.aws.amazon.com/ja_jp/AmazonS3/latest/userguide/access-control-block-public-access.html

Question: 160

A solutions architect needs to assess a newly acquired company's portfolio of applications and databases. The solutions architect must create a business case to migrate the portfolio to AWS. The newly acquired company runs applications in an on-premises data center. The data center is not well documented. The solutions architect cannot immediately determine how many applications and databases exist. Traffic for the applications is variable. Some applications are batch processes that run at the end of each month.

The solutions architect must gain a better understanding of the portfolio before a migration to AWS can begin.

Which solution will meet these requirements?

- A. Use AWS Server Migration Service (AWS SMS) and AWS Database Migration Service (AWS DMS) to evaluate migration. Use AWS Service Catalog to understand application and database dependencies.
- B. Use AWS Application Migration Service. Run agents on the on-premises infrastructure. Manage the agents by using AWS Migration Hub. Use AWS Storage Gateway to assess local storage needs and database dependencies.
- C. Use Migration Evaluator to generate a list of servers. Build a report for a business case. Use AWS Migration Hub to view the portfolio. Use AWS Application Discovery Service to gain an understanding of application dependencies.
- D. Use AWS Control Tower in the destination account to generate an application portfolio. Use AWS Server Migration Service (AWS SMS) to generate deeper reports and a business case. Use a landing zone for core accounts and resources.

Answer: C

Explanation:

The company should use Migration Evaluator to generate a list of servers and build a report for a business case. The company should use AWS Migration Hub to view the portfolio and use AWS Application Discovery Service to gain an understanding of application dependencies. This solution will meet the requirements because Migration Evaluator is a migration assessment service that helps create a data-driven business case for AWS cloud planning and migration. Migration Evaluator provides a clear baseline of what the company is running today and projects AWS costs based on measured on-premises provisioning and utilization¹. Migration Evaluator can use an agentless collector to conduct broad-based discovery or securely upload exports from existing inventory tools². Migration Evaluator integrates with AWS Migration Hub, which is a service that provides a single location to track the progress of application migrations across multiple AWS and partner solutions³. Migration Hub also supports AWS Application Discovery Service, which is a service that helps systems integrators quickly and reliably plan application migration projects by automatically identifying

applications running in on-premises data centers, their associated dependencies, and their performance profile.

<https://aws.amazon.com/migration-evaluator/> <https://docs.aws.amazon.com/migration-evaluator/latest/userguide/what-is.html><https://aws.amazon.com/migration-hub/>
<https://aws.amazon.com/application-discovery/> <https://aws.amazon.com/server-migration-service/>
<https://aws.amazon.com/dms/> <https://docs.aws.amazon.com/controltower/latest/userguide/what-is-control-tower.html><https://aws.amazon.com/application-migration-service/>
<https://aws.amazon.com/storagegateway/>

Question: 161

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 in the 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:

In the past, EBS can be attached only to one ec2 instance but not anymore but there are limitations like - it works only on io1/io2 instance types and many others as described here.
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html> EFS has shareable storage. In terms of performance, Amazon EFS is optimized for workloads that require high levels of aggregate throughput and IOPS, whereas EBS is optimized for low-latency, random access I/O operations. Amazon EFS is designed to scale throughput and capacity automatically as your storage needs grow, while EBS volumes can be resized on demand.

Question: 162

A company runs a customer service center that accepts calls and automatically sends all customers a

managed, interactive, two-way experience survey by text message.

The applications that support the customer service center run on machines that the company hosts in an on-premises data center. The hardware that the company uses is old, and the company is experiencing downtime with the system. The company wants to migrate the system to AWS to improve reliability.

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

- A. Use Amazon Connect to replace the old call center hardware. Use Amazon Pinpoint to send text message surveys to customers.
- B. Use Amazon Connect to replace the old call center hardware. Use Amazon Simple Notification Service (Amazon SNS) to send text message surveys to customers.
- C. Migrate the call center software to Amazon EC2 instances that are in an Auto Scaling group. Use the EC2 instances to send text message surveys to customers.
- D. Use Amazon Pinpoint to replace the old call center hardware and to send text message surveys to customers.

Answer: A

Explanation:

Amazon Connect is a cloud-based contact center service that allows you to set up a virtual call center for your business. It provides an easy-to-use interface for managing customer interactions through voice and chat. Amazon Connect integrates with other AWS services, such as Amazon S3 and Amazon Kinesis, to help you collect, store, and analyze customer data for insights into customer behavior and trends. On the other hand, Amazon Pinpoint is a marketing automation and analytics service that allows you to engage with your customers across different channels, such as email, SMS, push notifications, and voice. It helps you create personalized campaigns based on user behavior and enables you to track user engagement and retention. While both services allow you to communicate with your customers, they serve different purposes. Amazon Connect is focused on customer support and service, while Amazon Pinpoint is focused on marketing and engagement.

Question: 163

A company is building a call center by using Amazon Connect. The company's operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers.

Which solution will provide DR with the LOWEST RTO?

- A. Create an AWS Lambda function to check the availability of the Amazon Connect instance and to send a notification to the operations team in case of unavailability. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes. After notification, instruct the operations team to use the AWS Management Console to provision a new Amazon Connect instance in a second Region. Deploy the contact flows, users, and claimed phone numbers by using an AWS CloudFormation template.
- B. Provision a new Amazon Connect instance with all existing users in a second Region. Create an AWS Lambda function to check the availability of the Amazon Connect instance. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes. In the event of an issue, configure the Lambda function to deploy an AWS

CloudFormation template that provisions contact flows and claimed numbers in the second Region.

C. Provision a new Amazon Connect instance with all existing contact flows and claimed phone numbers in a second Region. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance. Create an Amazon CloudWatch alarm for failed health checks. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions all users. Configure the alarm to invoke the Lambda function.

D. Provision a new Amazon Connect instance with all existing users and contact flows in a second Region. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance. Create an Amazon CloudWatch alarm for failed health checks. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers. Configure the alarm to invoke the Lambda function.

Answer: D

Explanation:

Option D provisions a new Amazon Connect instance with all existing users and contact flows in a second Region. It also sets up an Amazon Route 53 health check for the URL of the Amazon Connect instance, an Amazon CloudWatch alarm for failed health checks, and an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers. This option allows for the fastest recovery time because all the necessary components are already provisioned and ready to go in the second Region. In the event of a disaster, the failed health check will trigger the AWS Lambda function to deploy the CloudFormation template to provision the claimed phone numbers, which is the only missing component.

Question: 164

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:

Amazon Route 53 Resolver is a managed DNS resolver service from Route 53 that helps to create conditional forwarding rules to redirect query traffic¹. By associating the private hosted zone to all the VPCs, the solutions architect can enable DNS resolution for cloud.example.com within the VPCs. By creating a Route 53 inbound resolver in the shared services VPC, the solutions architect can enable DNS resolution for cloud.example.com from on-premises systems. By attaching all VPCs to the transit gateway, the solutions architect can enable connectivity between the VPCs and the on-premises network through AWS Direct Connect. By creating forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver, the solutions architect can direct DNS queries for cloud.example.com to the Route 53 Resolver endpoint in AWS. This solution will provide the highest performance as it leverages Route 53 Resolver's optimized routing and caching capabilities.

Reference: 1: <https://aws.amazon.com/route53/resolver/>

Question: 165

A company is providing weather data over a REST-based API to several customers. The API is hosted by Amazon API Gateway and is integrated with different AWS Lambda functions for each API operation. The company uses Amazon Route 53 for DNS and has created a resource record of weather.example.com. The company stores data for the API in Amazon DynamoDB tables. The company needs a solution that will give the API the ability to fail over to a different AWS Region. Which solution will meet these requirements?

- A. Deploy a new set of Lambda functions in a new Region. Update the API Gateway API to use an edge-optimized API endpoint with Lambda functions from both Regions as targets. Convert the DynamoDB tables to global tables.
- B. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a multivalued answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.
- C. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a failover record. Enable target health monitoring. Convert the DynamoDB tables to global tables.
- D. Deploy a new API Gateway API in a new Region. Change the Lambda functions to global functions. Change the Route 53 DNS record to a multivalued answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.

Answer: C

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/dns-failover.html>

Question: 166

A company uses AWS Organizations with a single OU named Production to manage multiple accounts. All

accounts are members of the Production OU Administrators use deny list SCPs in the root of the organization to manage access to restricted services.

The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization. Once onboarded, the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.

Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

- A. Remove the organization's root SCPs that limit access to AWS Config. Create AWS Service Catalog products for the company's standard AWS Config rules and deploy them throughout the organization, including the new account.
- B. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config actions. Move the new account to the Production OU when adjustments to AWS Config are complete.
- C. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only. Temporarily apply an SCP to the organization's root that allows AWS Config actions for principals only in the new account.
- D. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config actions. Move the organization's root SCP to the Production OU. Move the new account to the Production OU when adjustments to AWS Config are complete.

Answer: D

Explanation:

An SCP at a lower level can't add a permission after it is blocked by an SCP at a higher level. SCPs can only filter; they never add permissions. So you need to create a new OU for the new account, assign an SCP, and move the root SCP to Production OU. Then move the new account to production OU when AWS config is done.

Question: 167

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

Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

- A. Enable Aurora Auto Scaling for Aurora Replicas. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.
- B. Enable Aurora Auto Scaling for Aurora writers. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.

- C. Enable Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing and sticky sessions enabled.
- D. Enable Aurora Scaling for Aurora writers. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

Explanation:

Aurora Auto Scaling enables your Aurora DB cluster to handle sudden increases in connectivity or workload. When the connectivity or workload decreases, Aurora AutoScaling removes unnecessary Aurora Replicas so that you don't pay for unused provisioned DB instances

Question: 168

A company runs an application on AWS. The company curates data from several different sources. The company uses proprietary algorithms to perform data transformations and aggregations. After the company performs ETL processes, the company stores the results in Amazon Redshift tables. The company sells this data to other companies. The company downloads the data as files from the Amazon Redshift tables and transmits the files to several data customers by using FTP. The number of data customers has grown significantly. Management of the data customers has become difficult.

The company will use AWS Data Exchange to create a data product that the company can use to share data with customers. The company wants to confirm the identities of the customers before the company shares data. The customers also need access to the most recent data when the company publishes the data.

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

- A. Use AWS Data Exchange for APIs to share data with customers. Configure subscription verification. In the AWS account of the company that produces the data, create an Amazon API Gateway Data API service integration with Amazon Redshift. Require the data customers to subscribe to the data product.
- B. In the AWS account of the company that produces the data, create an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster. Configure subscription verification. Require the data customers to subscribe to the data product.
- C. Download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically. Use AWS Data Exchange for S3 to share data with customers. Configure subscription verification. Require the data customers to subscribe to the data product.
- D. Publish the Amazon Redshift data to an Open Data on AWS Data Exchange. Require the customers to subscribe to the data product in AWS Data Exchange. In the AWS account of the company that produces the data, attach IAM resource-based policies to the Amazon Redshift tables to allow access only to verified AWS accounts.

Answer: C

Explanation:

The company should download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically and use AWS Data Exchange for S3 to share data with customers. The company should configure subscription verification and require the data customers to subscribe to the data product. This solution will meet the requirements with the least operational overhead because AWS Data Exchange for S3 is a feature that enables data subscribers to access third-party data files directly from data providers' Amazon S3 buckets. Subscribers can easily use these files for their data analysis with AWS services without needing to create or manage data copies. Data providers can easily set up AWS Data Exchange for S3 on top of their existing S3 buckets to share direct access to an entire S3 bucket or specific prefixes and S3 objects. AWS Data Exchange automatically manages subscriptions, entitlements, billing, and payment¹.

The other options are not correct because:

Using AWS Data Exchange for APIs to share data with customers would not work because AWS Data Exchange for APIs is a feature that enables data subscribers to access third-party APIs directly from data providers' AWS accounts. Subscribers can easily use these APIs for their data analysis with AWS services without needing to manage API keys or tokens. Data providers can easily set up AWS Data Exchange for APIs on top of their existing API Gateway resources to share direct access to an entire API or specific routes and stages². However, this feature is not suitable for sharing data from Amazon Redshift tables, which are not exposed as APIs. Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not work because the Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client³. It is useful for building applications that interact with Amazon Redshift, but not for sharing data files with customers. Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not work because AWS Data Exchange does not support datashares for Amazon Redshift clusters. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data⁴. It is useful for sharing query results and views with other users, but not for sharing data files with customers.

Publishing the Amazon Redshift data to an Open Data on AWS Data Exchange would not work because Open Data on AWS Data Exchange is a feature that enables you to find and use free and public datasets from AWS customers and partners. It is useful for accessing open and free data, but not for confirming the identities of the customers or charging them for the data.

Reference: <https://aws.amazon.com/data-exchange/why-aws-data-exchange/s3/>
<https://aws.amazon.com/data-exchange/why-aws-data-exchange/api/>
<https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html>
<https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html>
<https://aws.amazon.com/data-exchange/open-data/>

Question: 169

A company consists of two separate business units. Each business unit has its own AWS account within a single organization in AWS Organizations. The business units regularly share sensitive documents with each other. To facilitate sharing, the company created an Amazon S3 bucket in each account and configured two-way replication between the S3 buckets. The S3 buckets have millions of objects.

Recently, a security audit identified that neither S3 bucket has encryption at rest enabled. Company policy requires that all documents must be stored with encryption at rest. The company wants to implement server-side encryption with Amazon S3 managed encryption keys (SSE-S3).

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

A. Turn on SSE-S3 on both S3 buckets. Use S3 Batch Operations to copy and encrypt the objects in the same location.

- B. Create an AWS Key Management Service (AWS KMS) key in each account. Turn on server-side encryption with AWS KMS keys (SSE-KMS) on each S3 bucket by using the corresponding KMS key in that AWS account. Encrypt the existing objects by using an S3 copy command in the AWS CLI.
- C. Turn on SSE-S3 on both S3 buckets. Encrypt the existing objects by using an S3 copy command in the AWS CLI.
- D. Create an AWS Key Management Service (AWS KMS) key in each account. Turn on server-side encryption with AWS KMS keys (SSE-KMS) on each S3 bucket by using the corresponding KMS key in that AWS account. Use S3 Batch Operations to copy the objects into the same location.

Answer: A

Explanation:

"The S3 buckets have millions of objects" If there are million of objects then you should use Batch operations.
<https://aws.amazon.com/blogs/storage/encrypting-objects-with-amazon-s3-batch-operations/>

Question: 170

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

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

Which solution will meet these requirements?

- A. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. Configure a rule that has a low numeric value that allows requests for domains in the allowed list. Associate the rule group with the VPC.
- B. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a Route 53 outbound endpoint. Associate the outbound endpoint with the VPC. Associate the domain list with the outbound endpoint.
- C. Create an Amazon Route 53 traffic flow policy to match the allowed domains. Configure the traffic flow policy to forward requests that match to the Route 53 Resolver. Associate the traffic flow policy with the VPC.
- D. Create an Amazon Route 53 outbound endpoint. Associate the outbound endpoint with the VPC. Configure a Route 53 traffic flow policy to forward requests for allowed domains to the outbound endpoint. Associate the traffic flow policy with the VPC.

Answer: A

Explanation:

The company should create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. The company should configure a DNS Firewall rule group with a rule that has a high numeric value,

that blocks all requests. The company should configure a rule that has a low numeric value that allows requests for domains in the allowed list. The company should associate the rule group with the VPC. This solution will meet the requirements because Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can create reusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs. You can specify lists of domain names to allow or block, and you can customize the responses for the DNS queries that you block¹. By creating a domain list with the allowed domains and a rule group with rules to allow or block requests based on the domain list, the company can enforce its security policy and control access to sites.

The other options are not correct because:

Configuring a Route 53 outbound endpoint and associating it with the VPC would not help with filtering outbound DNS traffic. A Route 53 outbound endpoint is a resource that enables you to forward DNS queries from your VPC to your network over AWS Direct Connect or VPN connections². It does not provide any filtering capabilities.

Creating a Route 53 traffic flow policy to match the allowed domains would not help with filtering outbound DNS traffic. A Route 53 traffic flow policy is a resource that enables you to route traffic based on multiple criteria, such as endpoint health, geographic location, and latency³. It does not provide any filtering capabilities.

Creating a Gateway Load Balancer (GWLB) would not help with filtering outbound DNS traffic. A GWLB is a service that enables you to deploy, scale, and manage third-party virtual appliances such as firewalls, intrusion detection and prevention systems, and deep packet inspection systems in the cloud⁴. It does not provide any filtering capabilities.

Reference: <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-dns-firewall.html>
<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-outbound-endpoints.html>
<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/traffic-flow.html>
<https://docs.aws.amazon.com/elasticloadbalancing/latest/gateway/introduction.html>

Question: 171

An external audit of a company's serverless application reveals IAM policies that grant too many permissions. These policies are attached to the company's AWS Lambda execution roles. Hundreds of the company's Lambda functions have broad access permissions, such as full access to Amazon S3 buckets and Amazon DynamoDB tables. The company wants each function to have only the minimum permissions that the function needs to complete its task.

A solutions architect must determine which permissions each Lambda function needs.

What should the solutions architect do to meet this requirement with the LEAST amount of effort?

- A. Set up Amazon CodeGuru to profile the Lambda functions and search for AWS API calls. Create an inventory of the required API calls and resources for each Lambda function. Create new IAM access policies for each Lambda function. Review the new policies to ensure that they meet the company's business requirements.
- B. Turn on AWS CloudTrail logging for the AWS account. Use AWS Identity and Access Management Access Analyzer to generate IAM access policies based on the activity recorded in the CloudTrail log. Review the generated policies to ensure that they meet the company's business requirements.
- C. Turn on AWS CloudTrail logging for the AWS account. Create a script to parse the CloudTrail log, search for AWS API calls by Lambda execution role, and create a summary report. Review the report. Create IAM access policies that provide more restrictive permissions for each Lambda function.
- D. Turn on AWS CloudTrail logging for the AWS account. Export the CloudTrail logs to Amazon S3. Use Amazon EMR to process the CloudTrail logs in Amazon S3 and produce a report of API calls and resources used by each

execution role. Create a new IAM access policy for each role. Export the generated roles to an S3 bucket.

Review the generated policies to ensure that they meet the company's business requirements.

Answer: B

Explanation:

IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. IAM Access Analyzer identifies resources shared with external principals by using logic-based reasoning to analyze the resource-based policies in your AWS environment. <https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html>

Question: 172

A company is running an application in the AWS Cloud. The application collects and stores a large amount of unstructured data in an Amazon S3 bucket. The S3 bucket contains several terabytes of data and uses the S3 Standard storage class. The data increases in size by several gigabytes every day. The company needs to query and analyze the data

a. The company does not access data that is more than 1-year-old. However, the company must retain all the data indefinitely for compliance reasons.

Which solution will meet these requirements MOST cost-effectively?

A. Use S3 Select to query the data. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.

B. Use Amazon Redshift Spectrum to query the data. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.

C. Use an AWS Glue Data Catalog and Amazon Athena to query the data. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.

D. Use Amazon Redshift Spectrum to query the data. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Intelligent-Tiering.

Answer: C

Explanation:

Generally, unstructured data should be converted structured data before querying them. AWS Glue can do that. <https://docs.aws.amazon.com/glue/latest/dg/schema-relationalize.html> <https://docs.aws.amazon.com/athena/latest/ug/glue-athena.html>

Question: 173

A company's interactive web application uses an Amazon CloudFront distribution to serve images from an Amazon S3 bucket. Occasionally, third-party tools ingest corrupted images into the S3 bucket. This image corruption causes a poor user experience in the application later. The company has successfully implemented and tested Python logic to detect corrupt images.

A solutions architect must recommend a solution to integrate the detection logic with minimal latency between the ingestion and serving.

Which solution will meet these requirements?

- A. Use a Lambda@Edge function that is invoked by a viewer-response event.
- B. Use a Lambda@Edge function that is invoked by an origin-response event.
- C. Use an S3 event notification that invokes an AWS Lambda function.
- D. Use an S3 event notification that invokes an AWS Step Functions state machine.

Answer: B

Explanation:

This solution will allow the detection logic to be run as soon as the image is uploaded to the S3 bucket, before it is served to users via the CloudFront distribution. This way, the detection logic can quickly identify any corrupted images and prevent them from being served to users, minimizing latency between ingestion and serving.

Reference: AWS Lambda@Edge

documentation:<https://docs.aws.amazon.com/lambda/latest/dg/lambda-edge.html> You can use Lambda@Edge to run your code in response to CloudFront events, such as a viewer request, an origin request, a response, or an error.

Question: 174

A company has VPC flow logs enabled for its NAT gateway. The company is seeing Action = ACCEPT for inbound traffic that comes from public IP address 198.51.100.2 destined for a private Amazon EC2 instance.

A solutions architect must determine whether the traffic represents unsolicited inbound connections from the internet. The first two octets of the VPC CIDR block are 203.0.

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

- A. Open the AWS CloudTrail console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
- B. Open the Amazon CloudWatch console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
- C. Open the AWS CloudTrail console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
- D. Open the Amazon CloudWatch console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter

the sum of bytes transferred by the source address and the destination address.

Answer: D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/vpc-analyze-inbound-traffic-nat-gateway/> by Cloudxie says "select appropriate log"

Question: 175

A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.

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

- A. Migrate to Amazon CloudWatch dashboards. Recreate the dashboards to match the existing Grafana dashboards. Use automatic dashboards where possible.
- B. Create an Amazon Managed Grafana workspace. Configure a new Amazon CloudWatch data source. Export dashboards from the existing Grafana instance. Import the dashboards into the new workspace.
- C. Create an AMI that has Grafana pre-installed. Store the existing dashboards in Amazon Elastic File System (Amazon EFS). Create an Auto Scaling group that uses the new AMI. Set the Auto Scaling group's minimum, desired, and maximum number of instances to one. Create an Application Load Balancer that serves at least two Availability Zones.
- D. Configure AWS Backup to back up the EC2 instance that runs Grafana once each hour. Restore the EC2 instance from the most recent snapshot in an alternate Availability Zone when required.

Answer: C

Explanation:

By creating an AMI that has Grafana pre-installed and storing the existing dashboards in Amazon Elastic File System (Amazon EFS) it allows for faster and more efficient scaling, and by creating an Auto Scaling group that uses the new AMI and setting the Auto Scaling group's minimum, desired, and maximum number of instances to one and creating an Application Load Balancer that serves at least two Availability Zones, it ensures high availability and minimized downtime.

Question: 176

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

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

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

Answer: A,B

Explanation:

<https://medium.com/@it.melnichenko/invoke-a-lambda-across-multiple-aws-accounts-8c094b2e70be>

Question: 177

A company runs an IoT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MQTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named `iot.example.com`. The company uses Amazon Route 53 as its DNS service.

The company stores the data in Amazon DynamoDB.

On several occasions, the amount of data has overloaded the MQTT broker and has resulted in lost

sensor data.

a. The company must improve the reliability of the solution.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. Use the Auto Scaling group as the target for the ALB. Update the DNS record in Route 53 to an alias record. Point the alias record to the ALB. Use the MQTT broker to store the data.
- B. Set up AWS IoT Core to receive the sensor data. Create and configure a custom domain to connect to AWS IoT Core. Update the DNS record in Route 53 to point to the AWS IoT Core Data-ATS endpoint. Configure an AWS IoT rule to store the data.
- C. Create a Network Load Balancer (NLB). Set the MQTT broker as the target. Create an AWS Global Accelerator accelerator. Set the NLB as the endpoint for the accelerator. Update the DNS record in Route 53 to a multivalue answer record. Set the Global Accelerator IP addresses as values. Use the MQTT broker to store the data.
- D. Set up AWS IoT Greengrass to receive the sensor data. Update the DNS record in Route 53 to point to the AWS IoT Greengrass endpoint. Configure an AWS IoT rule to invoke an AWS Lambda function to store the data.

Answer: A

Explanation:

it describes a solution that uses an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. The ALB distributes incoming traffic across the instances in the Auto Scaling group and allows for automatic scaling based on incoming traffic. The use of an alias record in Route 53 allows for easy updates to the DNS record without changing the IP address. This solution improves the reliability of the MQTT broker by allowing it to automatically scale based on incoming traffic, reducing the likelihood of lost data due to broker overload.

Reference: <https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/>

<https://aws.amazon.com/autoscaling/>

<https://aws.amazon.com/route53/>

Question: 178

A company runs a processing engine in the AWS Cloud. The engine processes environmental data from logistics centers to calculate a sustainability index. The company has millions of devices in logistics centers that are spread across Europe. The devices send information to the processing engine through a RESTful API. The API experiences unpredictable bursts of traffic. The company must implement a solution to process all data that the devices send to the processing engine. Data loss is unacceptable. Which solution will meet these requirements?

A. Create an Application Load Balancer (ALB) for the RESTful API. Create an Amazon Simple Queue Service (Amazon SQS) queue. Create a listener and a target group for the ALB. Add the SQS queue as the target. Use a container that runs in Amazon Elastic Container Service (Amazon ECS) with the

Fargate launch type to process messages in the queue.

B. Create an Amazon API Gateway HTTP API that implements the RESTful API. Create an Amazon Simple Queue Service (Amazon SQS) queue. Create an API Gateway service integration with the SQS queue. Create an AWS Lambda function to process messages in the SQS queue.

C. Create an Amazon API Gateway REST API that implements the RESTful API. Create a fleet of Amazon EC2 instances in an Auto Scaling group. Create an API Gateway Auto Scaling group proxy integration. Use the EC2 instances to process incoming data.

D. Create an Amazon CloudFront distribution for the RESTful API. Create a data stream in Amazon Kinesis Data Streams. Set the data stream as the origin for the distribution. Create an AWS Lambda function to consume and process data in the data stream.

Answer: A

Explanation:

it will use the ALB to handle the unpredictable bursts of traffic and route it to the SQS queue. The SQS queue

will act as a buffer to store incoming data temporarily and the container running in Amazon ECS with the Fargate launch type will process messages in the queue. This approach will ensure that all data is processed and prevent data loss.

Question: 179

A solutions architect must provide a secure way for a team of cloud engineers to use the AWS CLI to upload objects into an Amazon S3 bucket. Each cloud engineer has an IAM user. IAM access keys and a virtual multi-factor authentication (MFA) device. The IAM users for the cloud engineers are in a group that is named S3-access. The cloud engineers must use MFA to perform any actions in Amazon S3.

Which solution will meet these requirements?

- A. Attach a policy to the S3 bucket to prompt the IAM user for an MFA code when the IAM user performs actions on the S3 bucket. Use IAM access keys with the AWS CLI to call Amazon S3.
- B. Update the trust policy for the S3-access group to require principals to use MFA when principals assume the group. Use IAM access keys with the AWS CLI to call Amazon S3.
- C. Attach a policy to the S3-access group to deny all S3 actions unless MFA is present. Use IAM access keys with the AWS CLI to call Amazon S3.
- D. Attach a policy to the S3-access group to deny all S3 actions unless MFA is present. Request temporary credentials from AWS Security Token Service (AWS STS). Attach the temporary credentials in a profile that Amazon S3 will reference when the user performs actions in Amazon S3.

Answer: D

Explanation:

The company should attach a policy to the S3-access group to deny all S3 actions unless MFA is present. The company should request temporary credentials from AWS Security Token Service (AWS STS). The company should attach the temporary credentials in a profile that Amazon S3 will reference when the user performs actions in Amazon S3. This solution will meet the requirements because AWS STS is a service that enables you to request temporary, limited-privilege credentials for IAM users or for users that you authenticate (federated users). You can use MFA with AWS STS to provide an extra layer of security when requesting temporary credentials¹. You can use the `sts get-session-token` AWS CLI command to request temporary credentials that include an MFA token². You can then use these credentials with the AWS CLI to access Amazon S3 resources. To do this, you need to attach a policy to the IAM group that denies all S3 actions unless MFA is present³. You also need to create a profile in the AWS CLI configuration file that references the temporary credentials.

The other options are not correct because:

Attaching a policy to the S3 bucket to prompt the IAM user for an MFA code when the IAM user performs actions on the S3 bucket would not work because policies attached to S3 buckets cannot enforce MFA authentication. Policies attached to S3 buckets are resource-based policies that define what actions can be performed on the bucket and by whom. They do not have any logic to prompt for an MFA code or verify it. Updating the trust policy for the S3-access group to require principals to use MFA when principals assume the group would not work because trust policies are used for roles, not groups. Trust policies are policies that define which principals can assume a role. They do not apply to groups, which are collections of IAM users that share permissions.

Creating an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains and

configuring a DNS Firewall rule group with rules to allow or block requests based on the domain list would not help with enforcing MFA authentication for Amazon S3 actions. Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can create reusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs. You can specify lists of domain names to allow or block, and you can customize the responses for the DNS queries that you block. This feature is useful for controlling access to sites and blocking DNS-level threats, but not for requiring MFA authentication. Reference:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp.htmlhttps://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_enable_cliapi.html

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_sample-policies.html<https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-dns-firewall.html>

Question: 180

A company wants to optimize AWS data-transfer costs and compute costs across developer accounts within the company's organization in AWS Organizations. Developers can configure VPCs and launch Amazon EC2 instances in a single AWS Region. The EC2 instances retrieve approximately 1 TB of data each day from Amazon S3. The developer activity leads to excessive monthly data-transfer charges and NAT gateway processing charges between EC2 instances and S3 buckets, along with high compute costs. The company wants to proactively enforce approved architectural patterns for any EC2 instance and VPC infrastructure that developers deploy within the AWS accounts. The company does not want this enforcement to negatively affect the speed at which the developers can perform their tasks. Which solution will meet these requirements MOST cost-effectively?

- A. Create SCPs to prevent developers from launching unapproved EC2 instance types. Provide the developers with an AWS CloudFormation template to deploy an approved VPC configuration with S3 interface endpoints. Scope the developers' IAM permissions so that the developers can launch VPC resources only with CloudFormation.
- B. Create a daily forecasted budget with AWS Budgets to monitor EC2 compute costs and S3 data-transfer costs across the developer accounts. When the forecasted cost is 75% of the actual budget cost, send an alert to the developer teams. If the actual budget cost is 100%, create a budget action to terminate the developers' EC2 instances and VPC infrastructure.
- C. Create an AWS Service Catalog portfolio that users can use to create an approved VPC configuration with S3 gateway endpoints and approved EC2 instances. Share the portfolio with the developer accounts. Configure an AWS Service Catalog launch constraint to use an approved IAM role. Scope the developers' IAM permissions to allow access only to AWS Service Catalog.
- D. Create and deploy AWS Config rules to monitor the compliance of EC2 and VPC resources in the developer AWS accounts. If developers launch unapproved EC2 instances or if developers create VPCs without S3 gateway endpoints, perform a remediation action to terminate the unapproved resources.

Answer: C

Explanation:

This solution allows developers to quickly launch resources using pre-approved configurations and instance types, while also ensuring that the resources launched comply with the company's architectural patterns. This

can help reduce data transfer and compute costs associated with the resources. Using AWS Service Catalog also allows the company to control access to the approved configurations and resources through the use of IAM roles, while also allowing developers to quickly provision resources without negatively affecting their ability to perform their tasks.

Reference:

AWS Service Catalog:<https://aws.amazon.com/service-catalog/>

AWS Service Catalog

Constraints:<https://docs.aws.amazon.com/servicecatalog/latest/adminguide/constraints.html> **AWS Service Catalog Launch**

Constraints:<https://docs.aws.amazon.com/servicecatalog/latest/adminguide/launch-constraints.html>

Question: 181

A company plans to migrate a three-tiered web application from an on-premises data center to AWS. The company developed the UI by using server-side JavaScript libraries. The business logic and API tier uses a Python-based web framework. The data tier runs on a MySQL database. The company custom built the application to meet business requirements. The company does not want to re-architect the application. The company needs a solution to replatform the application to AWS with the least possible amount of development. The solution needs to be highly available and must reduce operational overhead.

Which solution will meet these requirements?

- A. Deploy the UI to a static website on Amazon S3. Use Amazon CloudFront to deliver the website. Build the business logic in a Docker image. Store the image in Amazon Elastic Container Registry (Amazon ECR). Use Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type to host the website with an Application Load Balancer in front. Deploy the data layer to an Amazon Aurora MySQL DB cluster.
- B. Build the UI and business logic in Docker images. Store the images in Amazon Elastic Container Registry (Amazon ECR). Use Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type to host the UI and business logic applications with an Application Load Balancer in front. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB instance.
- C. Deploy the UI to a static website on Amazon S3. Use Amazon CloudFront to deliver the website. Convert the business logic to AWS Lambda functions. Integrate the functions with Amazon API Gateway. Deploy the data layer to an Amazon Aurora MySQL DB cluster.
- D. Build the UI and business logic in Docker images. Store the images in Amazon Elastic Container Registry (Amazon ECR). Use Amazon Elastic Kubernetes Service (Amazon EKS) with Fargate profiles to host the UI and business logic. Use AWS Database Migration Service (AWS DMS) to migrate the data layer to Amazon DynamoDB.

Answer: A

Explanation:

This solution utilizes Amazon S3 and CloudFront to deploy the UI as a static website, which can be done with minimal development effort. The business logic and API tier can be containerized in a Docker image and stored in Amazon Elastic Container Registry (ECR) and run on Amazon Elastic Container Service (ECS) with the Fargate launch type, which allows the application to be highly available with minimal operational overhead. The data

layer can be deployed on an Amazon Aurora MySQL DB cluster which is a fully managed relational database service.

Amazon Aurora provides high availability and performance for the data layer without the need for managing the underlying infrastructure.

Question: 182

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

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

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

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

Answer: A,E

Explanation:

A Direct Connect gateway allows you to connect multiple VPCs across different Regions to a Direct Connect connection¹. A public VIF allows you to access AWS public services such as EC2¹. A Site-to-Site VPN connection over the public VIF provides encryption and redundancy for the traffic between the on-premises data center and the VPCs². This solution is cheaper than setting up additional Direct Connect connections or using a private VIF with VPC peering.

Question: 183

A company ingests and processes streaming market data

a. The data rate is constant. A nightly process that calculates aggregate statistics is run, and each execution takes about 4 hours to complete. The statistical analysis is not mission critical to the business, and previous data points are picked up on the next execution if a particular run fails. The current architecture uses a pool of

Amazon EC2 Reserved Instances with 1-year reservations running full time to ingest and store the streaming data in attached Amazon EBS volumes. On-Demand EC2 instances are launched each night to perform the nightly processing, accessing the stored data from NFS shares on the ingestion servers, and terminating the nightly processing servers when complete. The Reserved Instance reservations are expiring, and the company needs to determine whether to purchase new reservations or implement a new design.

Which is the most cost-effective design?

A. Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon S3. Use a scheduled script to launch a fleet of EC2 On-Demand Instances each night to perform the batch processing of the S3 data. Configure the script to terminate the instances when the processing is complete.

B. Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon S3. Use AWS Batch with Spot Instances to perform nightly processing with a maximum Spot price that is 50% of the On-Demand price.

C. Update the ingestion process to use a fleet of EC2 Reserved Instances with 3-year reservations behind a Network Load Balancer. Use AWS Batch with Spot Instances to perform nightly processing with a maximum Spot price that is 50% of the On-Demand price.

D. Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon Redshift. Use Amazon EventBridge to schedule an AWS Lambda function to run nightly to query Amazon Redshift to generate the daily statistics.

Answer: B

Explanation:

Updating the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon S3 will reduce the need for EC2 instances and EBS volumes for data storage¹. Using AWS Batch with Spot Instances to perform nightly processing will leverage the cost savings of Spot Instances, which are up to 90% cheaper than On-Demand Instances². AWS Batch will also handle the scheduling and scaling of the processing jobs. Setting the maximum Spot price to 50% of the On-Demand price will reduce the chances of interruption and ensure that the processing is cost-effective.

Question: 184

A global manufacturing company plans to migrate the majority of its applications to AWS. However, the company is concerned about applications that need to remain within a specific country or in the company's central on-premises data center because of data regulatory requirements or requirements for latency of single-digit milliseconds. The company also is concerned about the applications that it hosts in some of its factory sites, where limited network infrastructure exists.

The company wants a consistent developer experience so that its developers can build applications once and deploy on premises, in the cloud, or in a hybrid architecture.

The developers must be able to use the same tools, APIs, and services that are familiar to them.

Which solution will provide a consistent hybrid experience to meet these requirements?

- A. Migrate all applications to the closest AWS Region that is compliant. Set up an AWS Direct Connect connection between the central on-premises data center and AWS. Deploy a Direct Connect gateway.
- B. Use AWS Snowball Edge Storage Optimized devices for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds. Retain the devices on premises.
- Deploy AWS Wavelength to host the workloads in the factory sites.
- C. Install AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds. Use AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites.
- D. Migrate the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds to an AWS Local Zone. Deploy AWS Wavelength to host the workloads in the factory sites.

Answer: C

Explanation:

Installing AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds will provide a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer premises¹. AWS Outposts allows customers to run some AWS services locally and connect to a broad range of services available in the local AWS Region¹. Using AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites will provide local compute and storage resources for locations with limited network infrastructure². AWS Snowball Edge devices can run Amazon EC2 instances and AWS Lambda functions locally and sync data with AWS when network connectivity is available².

Question: 185

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

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

- A. Configure AWS CloudTrail to log S3 data events.
- B. Configure S3 server access logging for the S3 bucket.
- C. Configure Amazon S3 to send object deletion events to Amazon Simple Email Service (Amazon SES).
- D. Configure Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic.
- E. Configure Amazon S3 to send the logs to Amazon Timestream with data storage tiering.
- F. Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.

Answer: A,D,F

Explanation:

Configuring AWS CloudTrail to log S3 data events will enable logging all activities for objects in the S3 bucket¹.

Data events are object-level API operations such as GetObject, DeleteObject, and PutObject¹. Configuring Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic will enable sending email notifications every time there is an attempt to delete data in the S3 bucket². EventBridge can route events from S3 to SNS, which can send emails to subscribers². Configuring a new S3 bucket to store the logs with an S3 Lifecycle policy will enable keeping the logs for 5 years in a cost-effective way³. A lifecycle policy can transition the logs to a cheaper storage class such as Glacier or delete them after a specified period of time³.

Question: 186

A solutions architect is reviewing a company's process for taking snapshots of Amazon RDS DB instances. The company takes automatic snapshots every day and retains the snapshots for 7 days.

The solutions architect needs to recommend a solution that takes snapshots every 6 hours and retains the snapshots for 30 days. The company uses AWS Organizations to manage all of its AWS accounts. The company needs a consolidated view of the health of the RDS snapshots.

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

- A. Turn on the cross-account management feature in AWS Backup. Create a backup plan that specifies the frequency and retention requirements. Add a tag to the DB instances. Apply the backup plan by using tags. Use AWS Backup to monitor the status of the backups.
- B. Turn on the cross-account management feature in Amazon RDS. Create a snapshot global policy that specifies the frequency and retention requirements. Use the RDS console in the management account to monitor the status of the backups.
- C. Turn on the cross-account management feature in AWS CloudFormation. From the management account, deploy a CloudFormation stack set that contains a backup plan from AWS Backup that specifies the frequency and retention requirements. Create an AWS Lambda function in the management account to monitor the status of the backups. Create an Amazon EventBridge rule in each account to run the Lambda function on a schedule.
- D. Configure AWS Backup in each account. Create an Amazon Data Lifecycle Manager lifecycle policy that specifies the frequency and retention requirements. Specify the DB instances as the target resource. Use the Amazon Data Lifecycle Manager console in each member account to monitor the status of the backups.

Answer: A

Explanation:

Turning on the cross-account management feature in AWS Backup will enable managing and monitoring backups across multiple AWS accounts that belong to the same organization in AWS Organizations¹. Creating a backup plan that specifies the frequency and retention requirements will enable taking snapshots every 6 hours and retaining them for 30 days². Adding a tag to the DB instances will enable applying the backup plan by using tags². Using AWS Backup to monitor the status of the backups will enable having a consolidated view of the health of the RDS snapshots¹.

Question: 187

A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes.

The company must maintain backups in a separate AWS Region. The company must be able to recover the EC2 instances and their configuration within 1 business day, with loss of no more than 1 day's worth of data.

a. The company has limited staff and needs a backup solution that optimizes operational efficiency and cost.

The company already has created an AWS CloudFormation template that can deploy the required network configuration in a secondary Region.

Which solution will meet these requirements?

A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbooks. Copy the snapshots to the secondary Region. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.

B. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volumes. In the event of a failure, launch the CloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.

C. Use AWS Backup to create a scheduled daily backup plan for the EC2 instances. Configure the backup task to copy the backups to a vault in the secondary Region. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.

D. Deploy EC2 instances of the same size and configuration to the secondary Region. Configure AWS DataSync daily to copy data from the primary Region to the secondary Region. In the event of a failure, launch the CloudFormation template and transfer usage to the secondary Region.

Answer: C

Explanation:

Using AWS Backup to create a scheduled daily backup plan for the EC2 instances will enable taking snapshots of the EC2 instances and their attached EBS volumes¹. Configuring the backup task to copy the backups to a vault in the secondary Region will enable maintaining backups in a separate Region¹. In the event of a failure, launching the CloudFormation template will enable deploying the network configuration in the secondary Region². Restoring the instance volumes and configurations from the backup vault will enable recovering the EC2 instances and their data¹. Transferring usage to the secondary Region will enable resuming operations².

Question: 188

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company's finance team

has a data processing application that uses AWS Lambda and Amazon DynamoDB. The company's marketing team wants to access the data that is stored in the DynamoDB table.

The DynamoDB table contains confidential data.

a. The marketing team can have access to only specific attributes of data in the DynamoDB table. The finance team and the marketing team have separate AWS accounts.

What should a solutions architect do to provide the marketing team with the appropriate access to the DynamoDB table?

- A. Create an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table. Attach the SCP to the OU of the finance team.
- B. Create an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes (fine-grained access control). Establish trust with the marketing team's account. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.
- C. Create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). Attach the policy to the DynamoDB table. In the marketing team's account, create an IAM role that has permissions to access the DynamoDB table in the finance team's account.
- D. Create an IAM role in the finance team's account to access the DynamoDB table. Use an IAM permissions boundary to limit the access to the specific attributes. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.

Answer: C

Explanation:

The company should create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). The company should attach the policy to the DynamoDB table. In the marketing team's account, the company should create an IAM role that has permissions to access the DynamoDB table in the finance team's account. This solution will meet the requirements because a resource-based IAM policy is a policy that you attach to an AWS resource (such as a DynamoDB table) to control who can access that resource and what actions they can perform on it. You can use IAM policy conditions to specify fine-grained access control for DynamoDB items and attributes. For example, you can allow or deny access to specific attributes of all items in a table by matching on attribute names¹. By creating a resource-based policy that allows access to only specific attributes of the DynamoDB table and attaching it to the table, the company can restrict access to confidential data. By creating an IAM role in the marketing team's account that has permissions to access the DynamoDB table in the finance team's account, the company can enable **CROSS-account access**.

The other options are not correct because:

Creating an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table would not work because SCPs are policies that you can use with AWS Organizations to manage permissions in your organization's accounts. SCPs do not grant permissions; instead, they specify the maximum permissions that identities in an account can have². SCPs cannot be used to specify fine-grained access control for DynamoDB items and attributes.

Creating an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes and establishing trust with the marketing team's account would not work because IAM roles are

identities that you can create in your account that have specific permissions. You can use an IAM role to delegate access to users, applications, or services that don't normally have access to your AWS resources³. However, creating an IAM role in the finance team's account would not restrict access to specific attributes of the DynamoDB table; it would only allow cross-account access. The company would still need a resource-based policy attached to the table to enforce fine-grained access control. Creating an IAM role in the finance team's account to access the DynamoDB table and using an IAM permissions boundary to limit the access to the specific attributes would not work because IAM permissions boundaries are policies that you use to delegate permissions management to other users. You can use permissions boundaries to limit the maximum permissions that an identity-based policy can grant to an IAM entity (user or role)⁴. Permissions boundaries cannot be used to specify fine-grained access control for DynamoDB items and attributes.

Reference: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/specifying-conditions.html>https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

Question: 189

A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the us-east-1 Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.

Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.

Which solution meets these requirements?

- A. Add an Amazon CloudFront distribution. Configure the ALB as the origin.
- B. Add an Amazon API Gateway edge-optimized API endpoint to expose the APIs. Configure the ALB as the target.
- C. Add an accelerator in AWS Global Accelerator. Configure the ALB as the origin.
- D. Deploy the APIs to two additional AWS Regions: eu-west-1 and ap-southeast-2. Add latency-based routing records in Amazon Route 53.

Answer: C

Explanation:

Adding an accelerator in AWS Global Accelerator will enable improving the performance of the APIs for local and global users¹. AWS Global Accelerator is a service that uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies¹. Configuring the ALB as the origin will enable connecting the accelerator to the ALB that exposes the APIs². AWS Global Accelerator supports non-standard REST methods such as LINK, UNLINK, LOCK, and UNLOCK³.

Question: 190

A company is implementing a serverless architecture by using AWS Lambda functions that need to access a Microsoft SQL Server DB instance on Amazon RDS. The company has separate environments for development and production, including a clone of the database system.

The company's developers are allowed to access the credentials for the development database. However, the credentials for the production database must be encrypted with a key that only members of the IT security team's IAM user group can access. This key must be rotated on a regular basis.

What should a solutions architect do in the production environment to meet these requirements?

- A. Store the database credentials in AWS Systems Manager Parameter Store by using a SecureString parameter that is encrypted by an AWS Key Management Service (AWS KMS) customer managed key. Attach a role to each Lambda function to provide access to the SecureString parameter. Restrict access to the SecureString parameter and the customer managed key so that only the IT security team can access the parameter and the key.
- B. Encrypt the database credentials by using the AWS Key Management Service (AWS KMS) default Lambda key. Store the credentials in the environment variables of each Lambda function. Load the credentials from the environment variables in the Lambda code. Restrict access to the KMS key so that only the IT security team can access the key.
- C. Store the database credentials in the environment variables of each Lambda function. Encrypt the environment variables by using an AWS Key Management Service (AWS KMS) customer managed key. Restrict access to the customer managed key so that only the IT security team can access the key.
- D. Store the database credentials in AWS Secrets Manager as a secret that is associated with an AWS Key Management Service (AWS KMS) customer managed key. Attach a role to each Lambda function to provide access to the secret. Restrict access to the secret and the customer managed key so that only the IT security team can access the secret and the key.

Answer: D

Explanation:

Storing the database credentials in AWS Secrets Manager as a secret that is associated with an AWS Key Management Service (AWS KMS) customer managed key will enable encrypting and managing the credentials securely¹. AWS Secrets Manager helps you to securely encrypt, store, and retrieve credentials for your databases and other services². Attaching a role to each Lambda function to provide access to the secret will enable retrieving the credentials programmatically¹. Restricting access to the secret and the customer managed key so that only members of the IT security team's IAM user group can access them will enable meeting the security requirements¹.

Question: 191

A company manufactures smart vehicles. The company uses a custom application to collect vehicle dat

a. The vehicles use the MQTT protocol to connect to the application.

The company processes the data in 5-minute intervals. The company then copies vehicle telematics data to on-premises storage. Custom applications analyze this data to detect anomalies.

The number of vehicles that send data grows constantly. Newer vehicles generate high volumes of data. The on-premises storage solution is not able to scale for peak traffic, which results in data loss. The company must modernize the solution and migrate the solution to AWS to resolve the scaling challenges.

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

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

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

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

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

Answer: B

Explanation:

Using AWS IoT Core to receive the vehicle data will enable connecting the smart vehicles to the cloud using the MQTT protocol¹. AWS IoT Core is a platform that enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, and enable applications to interact with devices even when they are offline². Configuring rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3 will enable processing and storing the vehicle data in a scalable and reliable way³. Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3. Creating an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies will enable analyzing the vehicle data using SQL queries or Apache Flink applications. Amazon Kinesis Data Analytics is a fully managed service that enables you to process and analyze streaming data using SQL or Java.

Question: 192

A company is designing a new website that hosts static content. The website will give users the ability to upload and download large files. According to company requirements, all data must be encrypted in transit and at rest. A solutions architect is building the solution by using Amazon S3 and Amazon CloudFront.

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

- A. Turn on S3 server-side encryption for the S3 bucket that the web application uses.
- B. Add a policy attribute of "aws:SecureTransport": "true" for read and write operations in the S3 ACLs.
- C. Create a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses.
- D. Configure encryption at rest on CloudFront by using server-side encryption with AWS KMS keys (SSE-KMS).
- E. Configure redirection of HTTP requests to HTTPS requests in CloudFront.
- F. Use the RequireSSL option in the creation of presigned URLs for the S3 bucket that the web application uses.

Answer: A,C,E

Explanation:

Turning on S3 server-side encryption for the S3 bucket that the web application uses will enable encrypting the data at rest using Amazon S3 managed keys (SSE-S3)¹. Creating a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses will enable enforcing encryption for all requests to the bucket². Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS³.

Question: 193

A company is using an organization in AWS Organizations to manage hundreds of AWS accounts. A solutions architect is working on a solution to provide baseline protection for the Open Web Application Security Project (OWASP) top 10 web application vulnerabilities. The solutions architect is using AWS WAF for all existing and new Amazon CloudFront distributions that are deployed within the organization.

Which combination of steps should the solutions architect take to provide the baseline protection? (Select THREE.)

- A. Enable AWS Config in all accounts.
- B. Enable Amazon GuardDuty in all accounts.
- C. Enable all features for the organization.
- D. Use AWS Firewall Manager to deploy AWS WAF rules in all accounts for all CloudFront distributions.
- E. Use AWS Shield Advanced to deploy AWS WAF rules in all accounts for all CloudFront distributions.
- F. Use AWS Security Hub to deploy AWS WAF rules in all accounts for all CloudFront distributions.

Answer: C,D,E

Explanation:

Enabling all features for the organization will enable using AWS Firewall Manager to centrally configure and manage firewall rules across multiple AWS accounts¹. Using AWS Firewall Manager to deploy AWS WAF rules in all accounts for all CloudFront distributions will enable providing baseline protection for the OWASP top 10 web application vulnerabilities². AWS Firewall Manager supports AWS WAF rules that can help protect against common web exploits such as SQL injection and crosssite scripting³. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS.

Question: 194

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

Which solution will meet these requirements?

- A. Use Amazon S3 for web hosting with Amazon API Gateway for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use Amazon Elastic Container Service (Amazon ECS) for business logic with Amazon SQS long polling for retaining failed orders.
- B. Use AWS Elastic Beanstalk for web hosting with Amazon API Gateway for database API services. Use Amazon MQ for order queuing. Use AWS Step Functions for business logic with Amazon S3 Glacier Deep Archive for retaining failed orders.
- C. Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
- D. Use Amazon Lightsail for web hosting with AWS AppSync for database API services. Use Amazon Simple Email Service (Amazon SES) for order queuing. Use Amazon Elastic Kubernetes Service (Amazon EKS) for business logic with Amazon OpenSearch Service for retaining failed orders.

Answer: C

Explanation:

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

This solution will allow you to:

- Host a static website on Amazon S3 without provisioning or managing servers¹.
- Use AWS AppSync to create a scalable GraphQL API that connects to your database and other data sources¹.
- Use Amazon SQS to decouple and scale your order processing microservices¹.
- Use AWS Lambda to run code for your business logic without provisioning or managing servers¹.
- Use an Amazon SQS dead-letter queue to retain messages that can't be processed by your Lambda function¹.

Question: 195

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

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

Which solution will meet these requirements MOST cost-effectively?

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

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

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

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

Answer: B

Explanation:

The company should use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. The company should create a cross-Region read replica for the DB instance. The company should set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. The company should run the EC2 instances at the minimum capacity in the DR Region. The company should use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. The company should increase the desired capacity of the Auto Scaling group. This solution will meet the requirements most cost-effectively because AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery. AWS DRS enables RPOs of seconds and RTOs of minutes¹. AWS DRS continuously replicates data from the source servers to a staging area subnet in the DR Region, where it uses low-cost storage and minimal compute resources to maintain ongoing replication. In the event of a disaster,

AWS DRS automatically converts the servers to boot and run natively on AWS and launches recovery instances on AWS within minutes². By using AWS DRS, the company can save costs by removing idle recovery site resources and paying for the full disaster recovery site only when needed. By creating a cross-Region read replica for the DB instance, the company can have a standby copy of its primary database in a different AWS Region³. By using infrastructure as code (IaC), the company can provision the new infrastructure in the DR Region in an automated and consistent way⁴. By using an Amazon Route 53 failover routing policy, the company can route traffic to a resource that is healthy or to another resource when the first resource becomes unavailable. The other options are not correct because:

Using AWS Backup to create cross-Region backups for the EC2 instances and the DB instance would not meet the RPO and RTO requirements. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. You can use AWS Backup to back up your application data across AWS services in your account and across accounts. However, AWS Backup does not provide continuous replication or fast recovery; it creates backups at scheduled intervals and requires manual restoration. Creating backups every 30 seconds would also incur high costs and network bandwidth.

Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data. Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data.

Reference: <https://aws.amazon.com/disaster-recovery/>

[https://docs.aws.amazon.com/drs/latest/userguide/what-is-](https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html)

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html#USER_ReadRepl

[.XRgn https://aws.amazon.com/cloudformation/](https://aws.amazon.com/cloudformation/)

[https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-](https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html)

[failover.html](https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html)<https://aws.amazon.com/backup/>

<https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html>[https://aws.amazon.com/data-](https://aws.amazon.com/data-exchange/)

[exchange/ https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html](https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html)

Question: 196

A company runs its sales reporting application in an AWS Region in the United States. The application uses an Amazon API Gateway Regional API and AWS Lambda functions to generate on-demand reports from data in an Amazon RDS for MySQL database. The frontend of the application is hosted on Amazon S3 and is accessed by users through an Amazon CloudFront distribution. The company is using Amazon Route 53 as the DNS service for the domain. Route 53 is configured with a simple routing policy to route traffic to the API Gateway API.

In the next 6 months, the company plans to expand operations to Europe. More than 90% of the database traffic is read-only traffic. The company has already deployed an API Gateway API and Lambda functions in the new Region.

A solutions architect must design a solution that minimizes latency for users who download reports.

Which solution will meet these requirements?

- A. Use an AWS Database Migration Service (AWS DMS) task with full load to replicate the primary database in the original Region to the database in the new Region. Change the Route 53 record to latency-based routing to connect to the API Gateway API.
- B. Use an AWS Database Migration Service (AWS DMS) task with full load plus change data capture (CDC) to replicate the primary database in the original Region to the database in the new Region. Change the Route 53 record to geolocation routing to connect to the API Gateway API.
- C. Configure a cross-Region read replica for the RDS database in the new Region. Change the Route 53 record to latency-based routing to connect to the API Gateway API.
- D. Configure a cross-Region read replica for the RDS database in the new Region. Change the Route 53 record to geolocation routing to connect to the API

Answer: C

Explanation:

The company should configure a cross-Region read replica for the RDS database in the new Region. The company should change the Route 53 record to latency-based routing to connect to the API Gateway API. This solution will meet the requirements because a cross-Region read replica is a feature that enables you to create a MariaDB, MySQL, Oracle, PostgreSQL, or SQL Server read replica in a different Region from the source DB instance. You can use cross-Region read replicas to improve availability and disaster recovery, scale out globally, or migrate an existing database to a new Region¹. By creating a cross-Region read replica for the RDS database in the new Region, the company can have a standby copy of its primary database that can serve read-only traffic from users in Europe. A latency-based routing policy is a feature that enables you to route traffic based on the latency between your users and your resources. You can use latency-based routing to route traffic to the resource that provides the best latency². By changing the Route 53 record to latency-based routing, the company can minimize latency for users who download reports by connecting them to the API Gateway API in the Region that provides the best response time.

The other options are not correct because:

Using AWS Database Migration Service (AWS DMS) to replicate the primary database in the original Region to the database in the new Region would not be as cost-effective or simple as using a cross-Region read replica. AWS DMS is a service that enables you to migrate relational databases, data warehouses, NoSQL databases, and other types of data stores. You can use AWS DMS to perform one-time migrations or continuous data replication with high availability and consolidate databases into a petabyte-scale data warehouse³. However, AWS DMS requires more configuration and management than creating a cross-Region read replica, which is fully managed by Amazon RDS. AWS DMS also incurs additional charges for replication instances and tasks.

Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery or minimizing latency. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data from an RDS database.

Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery or minimizing latency. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data from an RDS database.

Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS_Fea_Regions_DB-eng.Feature.CrossRegionReadReplicas.html<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-latency>
<https://aws.amazon.com/dms/> <https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html><https://aws.amazon.com/data-exchange/>
<https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html>

Question: 197

A company needs to optimize the cost of an AWS environment that contains multiple accounts in an organization in AWS Organizations. The company conducted cost optimization activities 3 years ago and purchased Amazon EC2 Standard Reserved Instances that recently expired.

The company needs EC2 instances for 3 more years. Additionally, the company has deployed a new serverless workload.

Which strategy will provide the company with the MOST cost savings?

- A. Purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. Purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs.
- B. Purchase a 1-year Compute Savings Plan with No Upfront payment in each member account. Use the Savings Plans recommendations in the AWS Cost Management console to choose the Compute Savings Plan.
- C. Purchase a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region. Purchase a 3-year Compute Savings Plan with No Upfront payment in the management account to cover any additional compute costs.
- D. Purchase a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account. Use the Savings Plans recommendations in the AWS Cost Management console to choose the EC2 Instance Savings Plan.

Answer: A

Explanation:

The company should purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. The company should purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs. This solution will provide the company with the most cost savings because Reserved Instances and Savings Plans are both pricing models that offer significant

discounts compared to On-Demand pricing. Reserved Instances are commitments to use a specific instance type and size in a single Region for a one- or three-year term. You can choose between three payment options: No Upfront, Partial Upfront, or All Upfront. The more you pay upfront, the greater the discount¹. Savings Plans are flexible pricing models that offer low prices on EC2 instances, Fargate, and Lambda usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one- or three-year term. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any EC2 instance regardless of Region, instance family, operating system, or tenancy, including those that are part of EMR, ECS, or EKS clusters, or launched by Fargate or Lambda. EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings². By purchasing the same Reserved Instances for an additional 3-year term with All Upfront payment, the company can lock in the lowest possible price for its EC2 instances that run continuously for 3 years. By purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account, the company can benefit from additional discounts on any other compute usage across its member accounts.

The other options are not correct because:

Purchasing a 1-year Compute Savings Plan with No Upfront payment in each member account would not provide as much cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. A 1-year term offers lower discounts than a 3-year term, and a No Upfront payment option offers lower discounts than an All Upfront payment option. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

Purchasing a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region would not provide as much cost savings as purchasing Reserved Instances for an additional 3-year term with All Upfront payment. An EC2 Instance Savings Plan offers lower discounts than Reserved Instances for the same instance family and Region. Also, a No Upfront payment option offers lower discounts than an All Upfront payment option.

Purchasing a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account would not provide as much flexibility or cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. An EC2 Instance Savings Plan applies only to a specific instance family within a Region and does not cover Fargate or Lambda usage. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

Reference: <https://aws.amazon.com/ec2/pricing/reserved-instances/>
<https://aws.amazon.com/savingsplans/>

Question: 198

A company is running an application in the AWS Cloud. The core business logic is running on a set of Amazon EC2 instances in an Auto Scaling group. An Application Load Balancer (ALB) distributes traffic to the EC2 instances. Amazon Route 53 record `api.example.com` is pointing to the ALB.

The company's development team makes major updates to the business logic. The company has a rule that when changes are deployed, only 10% of customers can receive the new logic during a testing window. A customer must use the same version of the business logic during the testing window.

How should the company deploy the updates to meet these requirements?

A. Create a second ALB, and deploy the new logic to a set of EC2 instances in a new Auto

Scaling group. Configure the ALB to distribute traffic to the EC2 instances. Update the Route 53 record to use weighted routing, and point the record to both of the ALBs.

B. Create a second target group that is referenced by the ALB. Deploy the new logic to EC2 instances in this new target group. Update the ALB listener rule to use weighted target groups. Configure ALB target group stickiness.

C. Create a new launch configuration for the Auto Scaling group. Specify the launch configuration to use the AutoScalingRollingUpdate policy, and set the MaxBatchSize option to 10. Replace the launch configuration on the Auto Scaling group. Deploy the changes.

D. Create a second Auto Scaling group that is referenced by the ALB. Deploy the new logic on a set of EC2 instances in this new Auto Scaling group. Change the ALB routing algorithm to least outstanding requests (LOR). Configure ALB session stickiness.

Answer: B

Explanation:

The company should create a second target group that is referenced by the ALB. The company should deploy the new logic to EC2 instances in this new target group. The company should update the ALB listener rule to use weighted target groups. The company should configure ALB target group stickiness. This solution will meet the requirements because weighted target groups are a feature that enables you to distribute traffic across multiple target groups using a single listener rule. You can specify a weight for each target group, which determines the percentage of requests that are routed to that target group. For example, if you specify two target groups, each with a weight of 10, each target group receives half the requests¹. By creating a second target group and deploying the new logic to EC2 instances in this new target group, the company can have two versions of its business logic running in parallel. By updating the ALB listener rule to use weighted target groups, the company can control how much traffic is sent to each version. By configuring ALB target group stickiness, the company can ensure that a customer uses the same version of the business logic during the testing window. Target group stickiness is a feature that enables you to bind a user's session to a specific target within a target group for the duration of the session².

The other options are not correct because:

Creating a second ALB and deploying the new logic to a set of EC2 instances in a new Auto Scaling group would not be as cost-effective or simple as using weighted target groups. A second ALB would incur additional charges and require more configuration and management. Updating the Route 53 record to use weighted routing would not ensure that a customer uses the same version of the business logic during the testing window, as DNS caching could affect how requests are routed. Creating a new launch configuration for the Auto Scaling group and replacing it on the Auto Scaling group would not allow for gradual traffic shifting between versions. A launch configuration is a template that an Auto Scaling group uses to launch EC2 instances. You can specify information such as the AMI ID, instance type, key pair, security groups, and block device mapping for your instances³. However, replacing the launch configuration on an Auto Scaling group would affect all instances in that group, not just 10% of customers.

Creating a second Auto Scaling group and changing the ALB routing algorithm to least outstanding requests (LOR) would not allow for controlled traffic shifting between versions. A second Auto Scaling group would require more configuration and management. The LOR routing algorithm is a feature that enables you to route traffic based on how quickly targets respond to requests. The load balancer selects a target from the target group with the fewest outstanding requests⁴. However, this algorithm does not take into account customer sessions or weights.

Reference: <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-listeners.html#listener-rules-weighted-target-groups>
<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/sticky-sessions.html>
<https://docs.aws.amazon.com/autoscaling/ec2/userguide/LaunchConfiguration.html>
<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html#routing-algorithm>

Question: 199

A company runs an application in an on-premises data center. The application gives users the ability to upload media files. The files persist in a file server. The web application has many users. The application server is overutilized, which causes data uploads to fail occasionally. The company frequently adds new storage to the file server. The company wants to resolve these challenges by migrating the application to AWS.

Users from across the United States and Canada access the application. Only authenticated users should have the ability to access the application to upload files. The company will consider a solution that refactors the application, and the company needs to accelerate application development.

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

- A. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances. Create an Auto Scaling group for the EC2 instances. Use an Application Load Balancer to distribute the requests. Modify the application to use Amazon S3 to persist the files. Use Amazon Cognito to authenticate users.
- B. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances. Create an Auto Scaling group for the EC2 instances. Use an Application Load Balancer to distribute the requests. Set up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application. Modify the application to use Amazon S3 to persist the files.
- C. Create a static website for uploads of media files. Store the static assets in Amazon S3. Use AWS AppSync to create an API. Use AWS Lambda resolvers to upload the media files to Amazon S3. Use Amazon Cognito to authenticate users.
- D. Use AWS Amplify to create a static website for uploads of media files. Use Amplify Hosting to serve the website through Amazon CloudFront. Use Amazon S3 to store the uploaded media files. Use Amazon Cognito to authenticate users.

Answer: D

Explanation:

The company should use AWS Amplify to create a static website for uploads of media files. The company should use Amplify Hosting to serve the website through Amazon CloudFront. The company should use Amazon S3 to store the uploaded media files. The company should use Amazon Cognito to authenticate users. This solution will meet the requirements with the least operational overhead because AWS Amplify is a complete solution that lets frontend web and mobile developers easily build, ship, and host full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as use cases evolve. No cloud expertise needed¹. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational

complexity and costs. AWS Amplify offers the following features and benefits: Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.

Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.

Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.

Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.

Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.

By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloudbackend in clicks. By using Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally. By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data². By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users³.

The other options are not correct because:

Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.

Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito. Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.

Reference: <https://aws.amazon.com/amplify/> <https://aws.amazon.com/s3/>
<https://aws.amazon.com/cognito/> <https://aws.amazon.com/mgn/>
<https://aws.amazon.com/appsync/> <https://aws.amazon.com/single-sign-on/>

Question: 200

A company has deployed its database on an Amazon RDS for MySQL DB instance in the us-east-1 Region. The company needs to make its data available to customers in Europe. The customers in Europe must have access to the same data as customers in the United States (US) and will not tolerate high application latency or stale data

a. The customers in Europe and the customers in the US need to write to the database. Both groups of customers need to see updates from the other group in real time.

Which solution will meet these requirements?

A. Create an Amazon Aurora MySQL replica of the RDS for MySQL DB instance. Pause application writes to the RDS DB instance. Promote the Aurora Replica to a standalone DB cluster. Reconfigure the application to use the Aurora database and resume writes. Add eu-west-1 as a secondary Region to the DB cluster. Enable write forwarding on the DB cluster. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.

B. Add a cross-Region replica in eu-west-1 for the RDS for MySQL DB instance. Configure the replica to replicate write queries back to the primary DB instance. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.

C. Copy the most recent snapshot from the RDS for MySQL DB instance to eu-west-1. Create a new RDS for MySQL DB instance in eu-west-1 from the snapshot. Configure MySQL logical replication from us-east-1 to eu-west-1. Enable write forwarding on the DB cluster. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.

D. Convert the RDS for MySQL DB instance to an Amazon Aurora MySQL DB cluster. Add eu-west-1 as a secondary Region to the DB cluster. Enable write forwarding on the DB cluster. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.

Answer: A

Explanation: The company should use AWS Amplify to create a static website for uploads of media files. The company should use Amplify Hosting to serve the website through Amazon CloudFront. The company should use Amazon S3 to store the uploaded media files. The company should use Amazon Cognito to authenticate users. This solution will meet the requirements with the least operational overhead because AWS Amplify is a complete solution that lets frontend web and mobile developers easily build, ship, and host full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as use cases evolve. No cloud expertise needed. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs. AWS Amplify offers the following features and benefits: Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.

Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.

Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.

Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.

Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.

By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloud backend in clicks. By using Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally.

By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data. By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users.

The other options are not correct because:

Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source

infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.

Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito. Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.

Reference:

<https://aws.amazon.com/amplify/>

<https://aws.amazon.com/s3/>

<https://aws.amazon.com/cognito/>

<https://aws.amazon.com/mgn/>

<https://aws.amazon.com/appsync/> <https://aws.amazon.com/single-sign-on/>

Question: 201

A solutions architect needs to improve an application that is hosted in the AWS Cloud. The application uses an Amazon Aurora MySQL DB instance that is experiencing overloaded connections. Most of the application's operations insert records into the database. The application currently stores credentials in a text-based configuration file.

The solutions architect needs to implement a solution so that the application can handle the current connection load. The solution must keep the credentials secure and must provide the ability to rotate the credentials automatically on a regular basis.

Which solution will meet these requirements?

- A. Deploy an Amazon RDS Proxy layer in front of the DB instance. Store the connection credentials as a secret in AWS Secrets Manager.
- B. Deploy an Amazon RDS Proxy layer in front of the DB instance. Store the connection credentials in AWS Systems Manager Parameter Store.
- C. Create an Aurora Replica. Store the connection credentials as a secret in AWS Secrets Manager.
- D. Create an Aurora Replica. Store the connection credentials in AWS Systems Manager Parameter Store.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html>

Question: 202

A company is deploying a new web-based application and needs a storage solution for the Linux application servers. The company wants to create a single location for updates to application data for all instances. The active dataset will be up to 100 GB in size. A solutions architect has determined that peak operations will occur for 3 hours daily and will require a total of 225 MiBps of read throughput.

The solutions architect must design a Multi-AZ solution that makes a copy of the data available in another AWS Region for disaster recovery (DR). The DR copy has an RPO of less than 1 hour.

Which solution will meet these requirements?

- A. Deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. Configure the file system for 75 MiBps of provisioned throughput. Implement replication to a file system in the DR Region.
- B. Deploy a new Amazon FSx for Lustre file system. Configure Bursting Throughput mode for the file system. Use AWS Backup to back up the file system to the DR Region.
- C. Deploy a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput. Enable Multi-Attach for the EBS volume. Use AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region.
- D. Deploy an Amazon FSx for OpenZFS file system in both the production Region and the DR Region. Create an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes.

Answer: A

Explanation:

The company should deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. The company should configure the file system for 75 MiBps of provisioned throughput. The company should implement replication to a file system in the DR Region. This solution will meet the requirements because Amazon EFS is a serverless, fully elastic file storage service that lets you share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files¹. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances. A Multi-AZ file system replicates data across multiple Availability Zones (AZs) within a Region, providing high availability and durability². By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput. Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system's size or burst credit balance³. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of

less than 1 hour.

The other options are not correct because:

Deploying a new Amazon FSx for Lustre file system would not provide a single location for updates to application data for all instances. Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance storage for compute workloads. However, it does not support concurrent write access from multiple instances. Using AWS Backup to back up the file system to the DR Region would not provide real-time replication of data. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. However, it does not support continuous data replication or cross-Region disaster recovery.

Deploying a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput would not provide a single location for updates to application data for all instances. Amazon EBS is a service that provides persistent block storage volumes for use with Amazon EC2 instances. However, it does not support concurrent access from multiple instances, unless Multi-Attach is enabled. Enabling Multi-Attach for the EBS volume would not provide Multi-AZ resilience or cross-Region replication. Multi-Attach is a feature that enables you to attach an EBS volume to multiple EC2 instances within the same Availability Zone. Using AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region would not provide real-time replication of data. AWS Elastic Disaster Recovery (AWS DRS) is a service that enables you to orchestrate and automate disaster recovery workflows across AWS Regions. However, it does not support continuous data replication or sub-hour RPOs.

Deploying an Amazon FSx for OpenZFS file system in both the production Region and the DR Region would not be as simple or cost-effective as using Amazon EFS. Amazon FSx for OpenZFS is a fully managed service that provides high-performance storage with strong data consistency and advanced data management features for Linux workloads. However, it requires more configuration and management than Amazon EFS, which is serverless and fully elastic. Creating an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes would not provide real-time replication of data. AWS DataSync is a service that enables you to transfer data between on-premises storage and AWS services, or between AWS services. However, it does not support continuous data replication or sub-minute RPOs.

Reference: <https://aws.amazon.com/efs/> <https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-azs><https://docs.aws.amazon.com/efs/latest/ug/performance.html#provisioned-throughput>

<https://docs.aws.amazon.com/efs/latest/ug/replication.html><https://aws.amazon.com/fsx/lustre/>

<https://aws.amazon.com/backup/> <https://aws.amazon.com/ebs/>

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html>

Question: 203

A company's public API runs as tasks on Amazon Elastic Container Service (Amazon ECS). The tasks run on AWS Fargate behind an Application Load Balancer (ALB) and are configured with Service Auto Scaling for the tasks based on CPU utilization. This service has been running well for several months.

Recently, API performance slowed down and made the application unusable. The company discovered that a significant number of SQL injection attacks had occurred against the API and that the API service had scaled to its maximum amount.

A solutions architect needs to implement a solution that prevents SQL injection attacks from reaching the ECS

API service. The solution must allow legitimate traffic through and must maximize operational efficiency.

Which solution meets these requirements?

- A. Create a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks.
- B. Create a new AWS WAF Bot Control implementation. Add a rule in the AWS WAF Bot Control managed rule group to monitor traffic and allow only legitimate traffic to the ALB in front of the ECS tasks.
- C. Create a new AWS WAF web ACL. Add a new rule that blocks requests that match the SQL database rule group. Set the web ACL to allow all other traffic that does not match those rules. Attach the web ACL to the ALB in front of the ECS tasks.
- D. Create a new AWS WAF web ACL. Create a new empty IP set in AWS WAF. Add a new rule to the web ACL to block requests that originate from IP addresses in the new IP set. Create an AWS Lambda function that scrapes the API logs for IP addresses that send SQL injection attacks, and add those IP addresses to the IP set. Attach the web ACL to the ALB in front of the ECS tasks.

Answer: C

Explanation:

The company should create a new AWS WAF web ACL. The company should add a new rule that blocks requests that match the SQL database rule group. The company should set the web ACL to allow all other traffic that does not match those rules. The company should attach the web ACL to the ALB in front of the ECS tasks. This solution will meet the requirements because AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications. You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked¹. By creating a new AWS WAF web ACL, the company can create a collection of rules that define the conditions for allowing or blocking web requests. By adding a new rule that blocks requests that match the SQL database rule group, the company can prevent SQL injection attacks from reaching the ECS API service. The SQL database rule group is a managed rule group provided by AWS that contains rules to protect against common SQL injection attack patterns². By setting the web ACL to allow all other traffic that does not match those rules, the company can ensure that legitimate traffic can access the API service. By attaching the web ACL to the ALB in front of the ECS tasks, the company can apply the web security rules to all requests that are forwarded by the load balancer.

The other options are not correct because:

Creating a new AWS WAF Bot Control implementation would not prevent SQL injection attacks from reaching the ECS API service. AWS WAF Bot Control is a feature that gives you visibility and control over common and pervasive bot traffic that can consume excess resources, skew metrics, cause downtime, or perform other undesired activities. However, it does not protect against SQL injection attacks, which are malicious attempts to execute unauthorized SQL statements against your database³.

Creating a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks would not prevent SQL injection attacks from reaching the ECS API service.

Monitoring mode is a feature that enables you to evaluate how your rules would perform without actually blocking any requests. However, this mode does not provide any protection against attacks, as it only logs and counts requests that match your rules⁴.

Creating a new AWS WAF web ACL and creating a new empty IP set in AWS WAF would not prevent SQL injection attacks from reaching the ECS API service. An IP set is a feature that enables you to specify a list of IP addresses or CIDR blocks that you want to allow or block based on their source IP address. However, this approach would not be effective or efficient against SQL injection attacks, as it would require constantly updating the IP set with new IP addresses of attackers, and it would not block attackers who use proxies or VPNs.

Reference:

<https://aws.amazon.com/waf/> <https://docs.aws.amazon.com/waf/latest/developerguide/aws-managed-rule-groups-list.html#sql-injection-rule-group>

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-bot-control.html>

<https://docs.aws.amazon.com/waf/latest/developerguide/web-acl-monitoring-mode.html>

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-ip-sets.html>

Question: 204

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. A public Application Load Balancer (ALB) will provide end users with access to the application. The company must prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack.

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

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

Answer: A,E

Explanation:

The company should create an Amazon CloudFront distribution with the ALB as the origin. The company should add a custom header and random value on the CloudFront domain. The company should configure the ALB to conditionally forward traffic if the header and value match. The company should also deploy an AWS WAF web ACL that includes an appropriate rule group. The company should associate the web ACL with the Amazon

CloudFront distribution. This solution will meet the requirements most cost-effectively because Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment¹. By creating an Amazon CloudFront distribution with the ALB as the origin, the company can improve the performance and availability of its application by caching static content at edge locations closer to end users. By adding a custom header and random value on the CloudFront domain, the company can prevent direct access to the ALB and ensure that only requests from CloudFront are forwarded to the ECS tasks. By configuring the ALB to conditionally forward traffic if the header and value match, the company can implement origin access identity (OAI) for its ALB origin. OAI is a feature that enables you to restrict access to your content by requiring users to access your content through CloudFront URLs². By deploying an AWS WAF web ACL that includes an appropriate rule group, the company can prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack. AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications. You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked³. By associating the web ACL with the Amazon CloudFront distribution, the company can apply the web security rules to all requests that are forwarded by CloudFront.

The other options are not correct because:

Deploying the application in two AWS Regions and configuring Amazon Route 53 to route to both Regions with equal weight would not prevent attacks or ensure business continuity. Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service that routes end users to Internet applications by translating names like `www.example.com` into numeric IP addresses⁴. However, routing traffic to multiple Regions would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and AWS WAF.

Configuring auto scaling for Amazon ECS tasks and creating a DynamoDB Accelerator (DAX) cluster would not prevent attacks or ensure business continuity. Auto scaling is a feature that enables you to automatically adjust your ECS tasks based on demand or a schedule. DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement. However, these features would not protect against attacks or provide failover in case of an outage. They would also increase operational complexity and costs compared to using CloudFront and AWS WAF.

Configuring Amazon ElastiCache to reduce overhead on DynamoDB would not prevent attacks or ensure business continuity. Amazon ElastiCache is a fully managed in-memory data store service that makes it easy to deploy, operate, and scale popular open-source compatible in-memory data stores. However, this service would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and AWS WAF. Reference:

<https://aws.amazon.com/cloudfront/>
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>
<https://aws.amazon.com/waf/>
<https://aws.amazon.com/route53/>
<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/service-auto-scaling.html>
<https://aws.amazon.com/dynamodb/dax/>
<https://aws.amazon.com/elasticache/>

Question: 205

A company needs to optimize the cost of backups for Amazon Elastic File System (Amazon EFS). A solutions architect has already configured a backup plan in AWS Backup for the EFS backups. The backup plan contains a rule with a lifecycle configuration to transition EFS backups to cold storage after 7 days and to keep the backups for an additional 90 days.

After 1 month, the company reviews its EFS storage costs and notices an increase in the EFS backup costs. The EFS backup cold storage produces almost double the cost of the EFS warm backup storage.

What should the solutions architect do to optimize the cost?

- A. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 1 day. Set the backup retention period to 30 days.
- B. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 8 days. Set the backup retention period to 30 days.
- C. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 1 day. Set the backup retention period to 90 days.
- D. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 8 days. Set the backup retention period to 98 days.

Answer: A

Explanation:

The cost of EFS backup cold storage is \$0.01 per GB-month, whereas the cost of EFS backup warm storage is \$0.05 per GB-month¹. Therefore, moving the backups to cold storage as soon as possible will reduce the storage cost. However, cold storage backups must be retained for a minimum of 90 days², otherwise they incur a pro-rated charge equal to the storage charge for the remaining days¹. Therefore, setting the backup retention period to 30 days will incur a penalty of 60 days of cold storage cost for each backup deleted. This penalty will still be lower than keeping the backups in warm storage for 7 days and then in cold storage for 83 days, which is the current configuration. Therefore, option A is the most cost-effective solution.

Question: 206

A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda

a. The application currently does not use

API keys to authorize requests. The API model is as follows:

GET/posts/[postid] to get post details

GET/users[userid] to get user details

GET/comments/[commentid] to get comments details

The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by marking the comments appears in real time. Which design should be used to reduce comment latency and improve user experience?

- A. Use edge-optimized API with Amazon CloudFront to cache API responses.
- B. Modify the blog application code to request GET comment[commented] every 10 seconds.
- C. Use AWS AppSync and leverage WebSockets to deliver comments.
- D. Change the concurrency limit of the Lambda functions to lower the API response time.

Answer: C

Explanation:

<https://docs.aws.amazon.com/appsync/latest/devguide/graphql-overview.html>

AWS AppSync is a fully managed GraphQL service that allows applications to securely access, manipulate, and receive data as well as real-time updates from multiple data sources¹. AWS AppSync supports GraphQL subscriptions to perform real-time operations and can push data to clients that choose to listen to specific events from the backend¹. AWS AppSync uses WebSockets to establish and maintain a secure connection between the clients and the API endpoint². Therefore, using AWS AppSync and leveraging WebSockets is a suitable design to reduce comment latency and improve user experience.

Question: 207

A solutions architect wants to cost-optimize and appropriately size Amazon EC2 instances in a single AWS account. The solutions architect wants to ensure that the instances are optimized based on CPU, memory, and network metrics.

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

- A. Purchase AWS Business Support or AWS Enterprise Support for the account.
- B. Turn on AWS Trusted Advisor and review any “Low Utilization Amazon EC2 Instances” recommendations.
- C. Install the Amazon CloudWatch agent and configure memory metric collection on the EC2 instances.
- D. Configure AWS Compute Optimizer in the AWS account to receive findings and optimization recommendations.
- E. Create an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems Of interest.

Answer: B,D

Explanation:

AWS Trusted Advisor is a service that provides real-time guidance to help users provision their resources following AWS best practices¹. One of the Trusted Advisor checks is “Low Utilization Amazon EC2 Instances”, which identifies EC2 instances that appear to be underutilized based on CPU, network I/O, and disk I/O metrics¹. This check can help users optimize the cost and size of their EC2 instances by recommending smaller or more appropriate instance types.

AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of AWS resources and generates optimization recommendations to reduce the cost and improve the performance of workloads². Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWS Fargate, and Lambda functions². For EC2 instances, Compute Optimizer evaluates the vCPUs, memory, storage, and other specifications, as well as the CPU utilization, network in and out, disk read and write, and other utilization metrics of currently running instances³. It then recommends optimal instance types based on price-performance tradeoffs.

Option A is incorrect because purchasing AWS Business Support or AWS Enterprise Support for the account will

not directly help with cost-optimization and sizing of EC2 instances. However, these support plans do provide access to more Trusted Advisor checks than the basic support plan¹. Option C is incorrect because installing the Amazon CloudWatch agent and configuring memory metric collection on the EC2 instances will not provide any optimization recommendations by itself. However, memory metrics can be used by Compute Optimizer to enhance its recommendations if enabled³.

Option E is incorrect because creating an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest will not help with cost-optimization and sizing of EC2 instances. Savings Plans are a flexible pricing model that offer lower prices on Amazon EC2 usage in exchange for a commitment to a consistent amount of usage for a 1- or 3-year term⁴. Savings Plans do not affect the configuration or utilization of EC2 instances.

Question: 208

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 policy to move all stored images to S3 Glacier Deep Archive after 6 months.

Answer: C

Explanation: S3 Event Notifications is a feature that allows users to receive notifications when certain events happen in an S3 bucket, such as object creation or deletion¹. Users can configure S3 Event Notifications to invoke an AWS Lambda function when a user stores an image in the bucket. Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources². The Lambda function can resize the image in place and store the original file in the same S3 bucket. This way, the solution can handle significant variance in demand and be reliable at enterprise scale. The solution can also rerun processing jobs in the event of failure by using the retry and dead-letter queue features of Lambda².

S3 Lifecycle is a feature that allows users to manage their objects so that they are stored cost-effectively throughout their lifecycle³. Users can create an S3 Lifecycle policy to move all stored images to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months. S3 Standard-IA is a storage class designed for data that is

accessed less frequently, but requires rapid access when needed⁴. It offers a lower storage cost than S3 Standard, but charges a retrieval fee. Therefore, moving the images to S3 Standard-IA after 6 months can reduce the storage cost for the solution.

Option A is incorrect because using AWS Step Functions to process the S3 event that occurs when a user stores an image is not necessary or cost-effective. AWS Step Functions is a service that lets users coordinate multiple AWS services into serverless workflows. However, for this use case, a single Lambda function can handle the image resizing task without needing Step Functions.

Option B is incorrect because using Amazon EventBridge to process the S3 event that occurs when a user uploads an image is not necessary or cost-effective. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. However, for this use case, S3 Event Notifications can directly invoke the Lambda function without needing EventBridge.

Option D is incorrect because using Amazon Simple Queue Service (Amazon SQS) to process the S3 event that occurs when a user stores an image is not necessary or cost-effective. Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices, distributed systems, and serverless applications. However, for this use case, S3 Event Notifications can directly invoke the Lambda function without needing SQS. Moreover, storing the resized file in an S3 bucket that uses S3 Standard-IA will incur a retrieval fee every time the file is accessed, which may not be cost-effective for frequently accessed files.

Question: 209

A company is using AWS Organizations to manage multiple AWS accounts. For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts.

A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of CloudFormation stacks. TrustedAccess has been enabled in Organizations.

What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

- A. Create a stack set in the Organizations member accounts. Use service-managed permissions. Set deployment options to deploy to an organization. Use CloudFormation StackSets drift detection.
- B. Create stacks in the Organizations member accounts. Use self-service permissions. Set deployment options to deploy to an organization. Enable the CloudFormation StackSets automatic deployment.
- C. Create a stack set in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets automatic deployment.
- D. Create stacks in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets drift detection.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-orgs-manage-auto-deployment.html>

Question: 210

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

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

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

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

Answer: C,D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-s3/>

Question: 211

A solutions architect is redesigning a three-tier application that a company hosts on-premises. The application provides personalized recommendations based on user profiles. The company already has an AWS account and has configured a VPC to host the application.

The frontend is a Java-based application that runs in on-premises VMs. The company hosts a personalization model on a physical application server and uses TensorFlow to implement the model. The personalization model uses artificial intelligence and machine learning (AI/ML). The company stores user information in a Microsoft SQL Server database. The web application calls the personalization model, which reads the user profiles from the database and provides recommendations.

The company wants to migrate the redesigned application to AWS.

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

- A. Use AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
- B. Export the personalization model. Store the model artifacts in Amazon S3. Deploy the model to Amazon SageMaker and create an endpoint. Host the Java application in AWS Elastic Beanstalk. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
- C. Use AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling groups. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to an EC2 instance.
- D. Containerize the personalization model and the Java application. Use Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy the model and the application to Amazon EKS. Host the node groups in a VPC. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to

Amazon RDS for SQL Server.

Answer: B

Explanation:

Amazon SageMaker is a fully managed machine learning service that allows users to build, train, and deploy machine learning models quickly and easily¹. Users can export their existing TensorFlow models and store the model artifacts in Amazon S3, a highly scalable and durable object storage service². Users can then deploy the model to Amazon SageMaker and create an endpoint that can be invoked by the web application to provide recommendations³. This way, the solution can leverage the AI/ML capabilities of Amazon SageMaker without having to rewrite the personalization model. AWS Elastic Beanstalk is a service that allows users to deploy and manage web applications without worrying about the infrastructure that runs those applications. Users can host their Java application in AWS Elastic Beanstalk and configure it to communicate with the Amazon SageMaker endpoint. This way, the solution can reduce the operational overhead of managing servers, load balancers, scaling, and application health monitoring.

AWS Database Migration Service (AWS DMS) is a service that helps users migrate databases to AWS quickly and securely. Users can use AWS DMS to migrate their SQL Server database to Amazon RDS for SQL Server, a fully managed relational database service that offers high availability, scalability, security, and compatibility. This way, the solution can reduce the operational overhead of managing database servers, backups, patches, and upgrades.

Option A is incorrect because using AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS is not cost-effective or scalable. AWS SMS is a service that helps users migrate on-premises workloads to AWS. However, for this use case, migrating the physical application server and the web application VMs to AWS will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS. Option C is incorrect because using AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling groups is not cost-effective or scalable. AWS Application Migration Service is a service that helps users migrate applications from on-premises or other clouds to AWS without making any changes to their applications. However, for this use case, migrating the personalization model and VMs to EC2 instances will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.

Option D is incorrect because containerizing the personalization model and the Java application and using Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy them to Amazon EKS is not necessary or cost-effective. Amazon EKS is a service that allows users to run Kubernetes on AWS without needing to install, operate, and maintain their own Kubernetes control plane or nodes. However, for this use case, containerizing and deploying the personalization model and the Java application will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk. Moreover, using S3 Glacier Deep Archive as a storage class for images will incur a high retrieval fee and latency for accessing them.

Question: 212

A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.

Which solution will meet these requirements?

- A. Create an alias for every new deployed version of the Lambda function. Use the AWS CLI `updatealias` command with the `routing-config` parameter to distribute the load.
- B. Deploy the application into a new CloudFormation stack. Use an Amazon Route 53 weighted routing policy to distribute the load.
- C. Create a version for every new deployed Lambda function. Use the AWS CLI `update-function-configuration` command with the `routing-config` parameter to distribute the load.
- D. Configure AWS CodeDeploy and use `CodeDeployDefault.OneAtATime` in the Deployment configuration to distribute the load.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-traffic-shifting/>

Question: 213

A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.

How can the company prevent users from accidentally deleting data in this way?

- A. Modify the CloudFormation templates to add a `DeletionPolicy` attribute to RDS and EBS resources.
- B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
- C. Modify IAM policies to deny deleting RDS and EBS resources that are tagged with an `"awsrcloudformation: stack-name"` tag.
- D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Answer: A

Explanation:

With the `DeletionPolicy` attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a `DeletionPolicy` attribute for each resource that you want to control. If a resource has no `DeletionPolicy` attribute, AWS CloudFormation deletes the resource by default. To keep a resource when its stack is deleted, specify `Retain` for that resource. You can use `retain` for any resource. For example, you can retain a nested stack, Amazon S3 bucket, or EC2 instance so that you can continue to use or modify those resources after you delete their stacks. <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

Question: 214

A company runs an intranet application on premises. The company wants to configure a cloud backup of the application. The company has selected AWS Elastic Disaster Recovery for this solution. The company requires that replication traffic does not travel through the public internet. The application also must not be accessible from the internet. The company does not want this solution to consume all available network bandwidth because other applications require bandwidth.

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

- A. Create a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway.
- B. Create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway.
- C. Create an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network.
- D. Create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network.
- E. During configuration of the replication servers, select the option to use private IP addresses for data replication.
- F. During configuration of the launch settings for the target servers, select the option to ensure that the Recovery instance's private IP address matches the source server's private IP address.

Answer: B,D,E

Explanation:

AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery¹. Users can set up AWS DRS on their source servers to initiate secure data replication to a staging area subnet in their AWS account, in the AWS Region they select. Users can then launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time.

To configure a cloud backup of the application with AWS DRS, users need to create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway. A VPC is a logically isolated section of the AWS Cloud where users can launch AWS resources in a virtual network that they define². A public subnet is a subnet that has a route to an internet gateway³. A virtual private gateway is the VPN concentrator on the Amazon side of the Site-to-Site VPN connection⁴. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in the VPC and the internet. Users need to create at least two public subnets for redundancy and high availability. Users need to create a virtual private gateway and attach it to the VPC to enable VPN connectivity between the on-premises network and the target AWS network. Users need to create an internet gateway and attach it to the VPC to enable internet access for the replication servers.

To ensure that replication traffic does not travel through the public internet, users need to create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network. AWS Direct Connect is a service that establishes a dedicated network connection from an on-premises network to one or more VPCs. A Direct Connect gateway is a globally available resource that allows users to connect multiple VPCs across different Regions to their on-premises networks using one or more Direct Connect connections. Users need to create an AWS Direct Connect connection between their on-premises network and an AWS Region. Users need to create a Direct Connect gateway and associate it with their VPC and their Direct Connect connection.

To ensure that the application is not accessible from the internet, users need to select the option to use private IP addresses for data replication during configuration of the replication servers. This option configures the replication servers with private IP addresses only, without assigning any public IP addresses or Elastic IP addresses. This way, the replication servers can only communicate with other resources within the VPC or through VPN connections.

Option A is incorrect because creating a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway is not necessary or cost-effective. A private subnet is a subnet that does not have a route to an internet gateway³. A NAT gateway is a highly available, managed Network Address Translation (NAT) service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances. Users do not need to create private subnets or NAT gateways for this use case, as they can use public subnets with private IP addresses for data replication.

Option C is incorrect because creating an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network will not ensure that replication traffic does not travel through the public internet. A Site-to-Site VPN connection consists of two VPN tunnels between an on-premises customer gateway device and a virtual private gateway in your VPC⁴. The VPN tunnels are encrypted using IPsec protocols, but they still use public IP addresses for communication. Users need to use AWS Direct Connect instead of Site-to-Site VPN for this use case.

Option F is incorrect because selecting the option to ensure that the Recovery instance's private IP address matches the source server's private IP address during configuration of the launch settings for the target servers will not ensure that the application is not accessible from the internet. This option configures the Recovery instance with an identical private IP address as its source server when launched in drills or recovery mode. However, this option does not prevent assigning public IP addresses or Elastic IP addresses to the Recovery instance. Users need to select the option to use private IP addresses for data replication instead.

Question: 215

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

Which solution will meet these requirements?

- A. Send event details to an Amazon Simple Notification Service (Amazon SNS) topic. Configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Add an on-failure destination to the function. Set an Amazon Simple Queue Service (Amazon SQS) queue as the target.
- B. Publish events to an Amazon Simple Queue Service (Amazon SQS) queue. Create an Amazon EC2 Auto Scaling group. Configure the Auto Scaling group to scale in and out based on the ApproximateAgeOfOldestMessage metric of the queue. Configure the application to write failed messages to a dead-letter queue.
- C. Write events to an Amazon DynamoDB table. Configure a DynamoDB stream for the table. Configure the stream to invoke an AWS Lambda function. Configure the Lambda function to process the events.
- D. Publish events to an Amazon EventBridge event bus. Create and run an application on an Amazon EC2 instance with an Auto Scaling group that is behind an Application Load Balancer (ALB). Set the ALB as the event bus target. Configure the event bus to retry events. Write messages to a dead-letter queue if the application cannot process the messages.

Answer: A

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fully managed pub/sub messaging service that enables users to send messages to multiple subscribers¹. Users can send event details to an Amazon SNS topic and configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources². Users can add an on-failure destination to the function and set an Amazon Simple Queue Service (Amazon SQS) queue as the target. Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices, distributed systems, and serverless applications³. This way, if a processing error occurs, the event will move into the separate queue for review.

Option B is incorrect because publishing events to an Amazon SQS queue and creating an Amazon EC2 Auto Scaling group will not have the ability to scale in and out based on the number of events that the solution receives. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Auto Scaling is a feature that helps users maintain application availability and allows them to scale their EC2 capacity up or down automatically according to conditions they define. However, for this use case, using SQS and EC2 will not take advantage of the serverless capabilities of Lambda and SNS.

Option C is incorrect because writing events to an Amazon DynamoDB table and configuring a DynamoDB stream for the table will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB Streams is a feature that captures data modification events in DynamoDB tables. Users can configure the stream to invoke a Lambda function, but they cannot configure an on-failure destination for the function. Option D is incorrect because publishing events to an Amazon EventBridge event bus and setting an Application Load Balancer (ALB) as the event bus target will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. An ALB is a load balancer that distributes incoming application traffic across multiple targets, such as EC2 instances, containers, IP addresses, Lambda functions, and virtual appliances. Users can configure EventBridge to retry events, but they cannot configure an on-failure destination for the ALB.

Question: 216

A company has IoT 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 IoT 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 **RealProvisioned Throughput Exceeded** error.

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

- A. Reshard the stream to increase the number of shards s 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:

<https://repost.aws/knowledge-center/kinesis-readprovisionedthroughputexceeded>

Follow Data Streams best practices

To mitigate ReadProvisionedThroughputExceeded exceptions, apply these best practices:

- Reshard your stream to increase the number of shards in the stream.
- Use consumers with enhanced fan-out. For more information about enhanced fan-out, see Developing custom consumers with dedicated throughput (enhanced fan-out).
- Use an error retry and exponential backoff mechanism in the consumer logic if

ReadProvisionedThroughputExceeded exceptions are encountered. For consumer applications that use an AWS SDK, the requests are retried by default.

Question: 217

An education company is running a web application used by college students around the world. The application runs in an Amazon Elastic Container Service (Amazon ECS) cluster in an Auto Scaling group behind an Application Load Balancer (ALB). A system administrator detected a weekly spike in the number of failed logic attempts. Which overwhelm the application's authentication service. All the failed login attempts originate from about 500 different IP addresses that change each week. A solutions architect must prevent the failed login attempts from overwhelming the authentication service.

Which solution meets these requirements with the MOST operational efficiency?

- A. Use AWS Firewall Manager to create a security group and security group policy to deny access from the IP addresses.
- B. Create an AWS WAF web ACL with a rate-based rule, and set the rule action to Block. Connect the web ACL to the ALB.
- C. Use AWS Firewall Manager to create a security group and security group policy to allow access only to specific CIDR ranges.
- D. Create an AWS WAF web ACL with an IP set match rule, and set the rule action to Block. Connect the web ACL to the ALB.

Answer: B

Explanation:

Question: 218

A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances.

To meet regulatory and business requirements, the company must make the following changes for data backups:

- * Backups must be retained based on custom daily, weekly, and monthly requirements.

* Backups must be replicated to at least one other AWS Region immediately after capture.
* The backup solution must provide a single source of backup status across the AWS environment.
* The backup solution must send immediate notifications upon failure of any resource backup. Which combination of steps will meet 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.

Answer: A,B,D

Explanation:

Cross region with AWS Backup:<https://docs.aws.amazon.com/aws-backup/latest/devguide/cross-region-backup.html>

Question: 219

A company has many separate AWS accounts and uses no central billing or management. Each AWS account hosts services for different departments in the company. The company has a Microsoft Azure Active Directory that is deployed.

A solution architect needs to centralize billing and management of the company's AWS accounts. The company wants to start using identify federation instead of manual user management. The company also wants to use temporary credentials instead of long-lived access keys.

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

- A. Create a new AWS account to serve as a management account. Deploy an organization in AWS Organizations. Invite each existing AWS account to join the organization. Ensure that each account accepts the invitation.
- B. Configure each AWS Account's email address to be aws+<account id>@example.com so that account management email messages and invoices are sent to the same place.
- C. Deploy AWS IAM Identity Center (AWS Single Sign-On) in the management account. Connect IAM Identity Center to the Azure Active Directory. Configure IAM Identity Center for automatic synchronization of users and groups.
- D. Deploy an AWS Managed Microsoft AD directory in the management account. Share the directory with all other accounts in the organization by using AWS Resource Access Manager (AWS RAM).
- E. Create AWS IAM Identity Center (AWS Single Sign-On) permission sets. Attach the permission sets to the appropriate IAM Identity Center groups and AWS accounts.
- F. Configure AWS Identity and Access Management (IAM) in each AWS account to use AWS Managed Microsoft AD for authentication and authorization.

Answer: A,C,E

Explanation:

Question: 220

A company processes environment data

a. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.

The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time.

Which solution will meet these requirements?

- A. Use Amazon Kinesis Data Firehouse to send the data to Amazon Redshift.
- B. Use Amazon Kinesis Data streams to send the data to Amazon DynamoDB.
- C. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to send the data to Amazon Aurora.
- D. Use Amazon Kinesis Data firehouse to send the data to Amazon Keyspaces (for Apache Cassandra).

Answer: B

Explanation:

Amazon Kinesis Data Streams is a service that enables real-time data ingestion and processing.

Amazon DynamoDB is a NoSQL database that does not require fixed schemas for storage. By using Kinesis Data Streams and DynamoDB, the company can send the JSON data to a database that can handle schemaless data in real time. Reference:

<https://docs.aws.amazon.com/streams/latest/dev/introduction.html><https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>

Question: 221

A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-through put. low-latency network connections between all to the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant.

Which solution will meet these requirements?

- A. Launch five new EC2 instances into a cluster placement group. Ensure that the EC2instance type supports enhanced networking.
- B. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zone. Attach an extra elastic network interface to each EC2 instance.
- C. Launch five new EC2 instances into a partition placement group. Ensure that the EC2 instance type supports enhanced networking.
- D. Launch five new EC2 instances into a spread placement group Attach an extra elastic network interface to each EC2 instance.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html#placement-groups-cluster>

Question: 222

A company has migrated a legacy application to the AWS Cloud. The application runs on three Amazon EC2 instances that are spread across three Availability Zones. One EC2 instance is in each Availability Zone. The EC2 instances are running in three private subnets of the VPC and are set up as targets for an Application Load Balancer (ALB) that is associated with three public subnets.

The application needs to communicate with on-premises systems. Only traffic from IP addresses in the company's IP address range are allowed to access the on-premises systems. The company's security team is bringing only one IP address from its internal IP address range to the cloud. The company has added this IP address to the allow list for the company firewall. The company also has created an Elastic IP address for this IP address.

A solutions architect needs to create a solution that gives the application the ability to communicate with the on-premises systems. The solution also must be able to mitigate failures automatically. Which solution will meet these requirements?

- A. Deploy three NAT gateways, one in each public subnet. Assign the Elastic IP address to the NAT gateways. Turn on health checks for the NAT gateways. If a NAT gateway fails a health check, recreate the NAT gateway and assign the Elastic IP address to the new NAT gateway.
- B. Replace the ALB with a Network Load Balancer (NLB). Assign the Elastic IP address to the NLB. Turn ON health checks for the NLB. In the case of a failed health check, redeploy the NLB in different subnets.
- C. Deploy a single NAT gateway in a public subnet. Assign the Elastic IP address to the NAT gateway. Use Amazon CloudWatch with a custom metric to monitor the NAT gateway. If the NAT gateway is unhealthy, invoke an AWS Lambda function to create a new NAT gateway in a different subnet. Assign the Elastic IP address to the new NAT gateway.
- D. Assign the Elastic IP address to the ALB. Create an Amazon Route 53 simple record with the Elastic IP address as the value. Create a Route 53 health check. In the case of a failed health check, recreate the ALB in different subnets.

Answer: C

Explanation:

to connect out from the private subnet you need a NAT gateway and since only one Elastic IP is whitelisted on the firewall, only one NAT gateway at a time and if an AZ failure happens, Lambda creates a new NAT gateway in a different AZ using the same Elastic IP. Don't be tempted to select D since an application that needs to connect to a private subnet whose outbound connections use the NAT gateway Elastic IP.

Question: 223

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

In each AWS account with a client, an interface endpoint has been created for the logging service and is available. The logging service running on EC2 instances with a Network Load Balancer (NLB) are deployed in different subnets. The clients are unable to submit logs using the VPC endpoint. Which combination of steps should a solutions architect take to resolve this issue? (Select TWO.)

A. Check that the NACL is attached to the logging service subnet to allow communications to and from the NLB subnets. Check that the NACL is attached to the NLB subnet to allow communications to and from the logging service subnets running on EC2 instances.

B. Check that the NACL is attached to the logging service subnets to allow communications to and from the interface endpoint subnets. Check that the NACL is attached to the interface endpoint subnet to allow communications to and from the logging service subnets running on EC2 instances.

C. Check the security group for the logging service running on the EC2 instances to ensure it allows

ingress from the NLB subnets.

D. Check the security group for the logging service running on EC2 instances to ensure it allows ingress from the clients.

E. Check the security group for the NLB to ensure it allows ingress from the interface endpoint subnets.

Answer: A,C

Explanation:

Question: 224

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

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

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

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

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

Answer: B

Explanation:

Transit GW + Direct Connect GW + Transit VIF + enabled SiteLink if two different DX locations
<https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-aws-direct-connect-sitelink/>

Question: 225

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

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

Which solution meets these requirements?

- A. Configure a Multi-AZ Auto Scaling group using the application's AMI. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Use Amazon Inspector to monitor traffic to the ALB and EC2 instances. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application.
- B. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Logs. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.
- C. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as targets. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.
- D. Configure a Multi-AZ Auto Scaling group using the application's AMI. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the WebACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

Answer: D

Explanation:

<https://docs.aws.amazon.com/waf/latest/developerguide/marketplace-managed-rule-groups.html>

Question: 226

A company has developed a hybrid solution between its data center and AWS. The company uses Amazon VPC and Amazon EC2 instances that send application logs to Amazon CloudWatch. The EC2 instances read data from multiple relational databases that are hosted on premises.

The company wants to monitor which EC2 instances are connected to the databases in near-real time. The company already has a monitoring solution that uses Splunk on premises. A solutions architect needs to determine how to send networking traffic to Splunk.

How should the solutions architect meet these requirements?

- A. Enable VPC flows logs, and send them to CloudWatch. Create an AWS Lambda function to periodically export the CloudWatch logs to an Amazon S3 bucket by using the pre-defined export function. Generate ACCESS_KEY and SECRET_KEY AWS credentials. Configure Splunk to pull the logs from the S3 bucket by using those credentials.
- B. Create an Amazon Kinesis Data Firehose delivery stream with Splunk as the destination. Configure a pre-processing AWS Lambda function with a Kinesis Data Firehose stream processor that extracts individual log events from records sent by CloudWatch Logs subscription filters. Enable VPC flows logs, and send them to CloudWatch. Create a CloudWatch Logs subscription that sends log events to the Kinesis Data Firehose delivery stream.
- C. Ask the company to log every request that is made to the databases along with the EC2 instance IP address. Export the CloudWatch logs to an Amazon S3 bucket. Use Amazon Athena to query the logs grouped by database name. Export Athena results to another S3 bucket. Invoke an AWS Lambda function to automatically send any new file that is put in the S3 bucket to Splunk.
- D. Send the CloudWatch logs to an Amazon Kinesis data stream with Amazon Kinesis Data Analytics for SQL Applications. Configure a 1 -minute sliding window to collect the events. Create a SQL query that uses the anomaly detection template to monitor any networking traffic anomalies in near-real time. Send the result to an Amazon Kinesis Data Firehose delivery stream with Splunk as the destination.

Answer: B

Explanation:

<https://docs.aws.amazon.com/firehose/latest/dev/creating-the-stream-to-splunk.html>

Question: 227

A company has built a high performance computing (HPC) cluster in AWS for a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.

Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)

- A. Ensure the HPC cluster is launched within a single Availability Zone.

- B. Launch the EC2 instances and attach elastic network interfaces in multiples of four.
- C. Select EC2 Instance types with an Elastic Fabric Adapter (EFA) enabled.
- D. Ensure the cluster is launched across multiple Availability Zones.
- E. Replace Amazon EFS with multiple Amazon EBS volumes in a RAID array.
- F. Replace Amazon EFS with Amazon FSx for Lustre.

Answer: A,C,F

Explanation:

A . High performance computing (HPC) workload cluster should be in a single AZ.

C . Elastic Fabric Adapter (EFA) is a network device that you can attach to your Amazon EC2 instances to accelerate High Performance Computing (HPC)

F . Amazon FSx for Lustre - Use it for workloads where speed matters, such as machine learning, high performance computing (HPC), video processing, and financial modeling.

Cluster - packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

Question: 228

A company has millions of objects in an Amazon S3 bucket. The objects are in the S3 Standard storage class. All the S3 objects are accessed frequently. The number of users and applications that access the objects is increasing rapidly. The objects are encrypted with server-side encryption with AWS KMS Keys (SSE-KMS).

A solutions architect reviews the company's monthly AWS invoice and notices that AWS KMS costs are increasing because of the high number of requests from Amazon S3. The solutions architect needs to optimize costs with minimal changes to the application.

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

- A. Create a new S3 bucket that has server-side encryption with customer-provided keys (SSE-C) as the encryption type. Copy the existing objects to the new S3 bucket. Specify SSE-C.
- B. Create a new S3 bucket that has server-side encryption with Amazon S3 managed keys (SSE-S3) as the encryption type. Use S3 Batch Operations to copy the existing objects to the new S3 bucket. Specify SSE-S3.
- C. Use AWS CloudHSM to store the encryption keys. Create a new S3 bucket. Use S3 Batch Operations to copy the existing objects to the new S3 bucket. Encrypt the objects by using the keys from CloudHSM.
- D. Use the S3 Intelligent-Tiering storage class for the S3 bucket. Create an S3 Intelligent-Tiering archive configuration to transition objects that are not accessed for 90 days to S3 Glacier Deep Archive.

Answer: B

Explanation:

To reduce the volume of Amazon S3 calls to AWS KMS, use Amazon S3 bucket keys, which are protected encryption keys that are reused for a limited time in Amazon S3. Bucket keys can reduce costs for AWS KMS requests by up to 99%. You can configure a bucket key for all objects in an Amazon S3 bucket, or for a specific object in an Amazon S3

bucket. https://docs.aws.amazon.com/fr_fr/kms/latest/developerguide/services-s3.html

Question: 229

A solutions architect is designing an AWS account structure for a company that consists of multiple teams. All the teams will work in the same AWS Region. The company needs a VPC that is connected to the on-premises network. The company expects less than 50 Mbps of total traffic to and from the on-premises network. Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

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

Answer: B,D

Explanation:

Question: 230

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 interAvailability Zone data transfer usage.

Answer: C,D

Explanation:

Cross-zone load balancing enables traffic to be distributed evenly across all registered instances in all enabled Availability Zones. However, this also increases data transfer charges between Availability Zones. By turning off cross-zone load balancing, the service provider applications can reduce interAvailability Zone data transfer costs. Similarly, by using the Availability Zone-specific endpoint service, the service consumer applications can ensure that they connect to the nearest service provider application in the same Availability Zone, avoiding cross-Availability Zone data transfer charges. Reference: <https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#cross-zone-load-balancing><https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html#vpce-interface-dns>

Question: 231

A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times for all task, primary, and core nodes. The EMR tasks run each morning, starting at 1:00 AM, and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.

The solutions architect must review the architecture and suggest a solution to minimize the compute costs. Which solution should the solutions architect recommend to meet these requirements?

- A. Launch all task, primary, and core nodes on Spot Instances in an instance fleet. Terminate the cluster, including all instances, when the processing is completed.
- B. Launch the primary and core nodes on On-Demand Instances. Launch the task nodes on Spot Instances in an instance fleet. Terminate the cluster, including all instances, when the processing is completed. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
- C. Continue to launch all nodes on On-Demand Instances. Terminate the cluster, including all instances, when the processing is completed. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
- D. Launch the primary and core nodes on On-Demand Instances. Launch the task nodes on Spot Instances in an instance fleet. Terminate only the task node instances when the processing is completed. Purchase Compute Savings Plans to cover the On-Demand Instance usage.

Answer: A

Explanation:

Amazon EC2 Spot Instances offer spare compute capacity at steep discounts compared to On-Demand prices. Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back. Amazon EMR can handle Spot interruptions gracefully by decommissioning the nodes and redistributing the tasks to other nodes. By launching all nodes on Spot Instances in an instance fleet, the solutions architect can minimize the compute costs of the EMR cluster. An instance fleet is a collection of EC2 instances with different types and sizes that EMR automatically provisions to meet a defined target capacity. By terminating the cluster when the processing is completed, the solutions architect can avoid paying for idle resources.

Reference:

<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-managed-scaling.html>

<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-instance-fleet.html>
<https://aws.amazon.com/blogs/big-data/optimizing-amazon-emr-for-resilience-and-cost-with-capacity-optimized-spot-instances/>

Question: 232

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:

<https://joe.blog.freemansoft.com/2020/04/protect-data-in-cloud-with-s3-access.html>

Question: 233

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 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,C,E

Explanation:

AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. By creating a management account and inviting all the existing accounts to join the organization, the solutions architect can track and consolidate expenditures for all the accounts using AWS Cost Management tools such as AWS Cost Explorer and AWS Budgets. An organizational unit (OU) is a group of accounts within an organization that can be used to apply policies and simplify management. A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. By creating an OU that includes all the development teams and applying an SCP that allows the creation of resources only in Regions that are in the United States, the solutions architect can ensure that the company meets its compliance requirements and avoids unwanted charges from other Regions. An IAM role is an identity with permission policies that determine what the identity can and cannot do in AWS. By creating an IAM role in the management account and allowing the finance team users to assume it, the solutions architect can give them

access to view the Billing and Cost Management console without sharing credentials or creating additional users. Reference:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_introduction.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
<https://docs.aws.amazon.com/aws-cost-management/latest/userguide/what-is-costmanagement.html>

Question: 234 A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.

The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images. When a new image version is uploaded, the new image version receives a unique tag. The company needs a solution that inspects new image versions for common vulnerabilities and exposures. The solution must automatically delete new image tags that have Critical or High severity findings. The solution also must notify the development team when such a deletion occurs. Which solution meets these requirements?

- A. Configure scan on push on the repository Use Amazon EventBridge to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical or High severity findings. Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS).
- B. Configure scan on push on the repository Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Lambda function when a new message is added to the SQS queue. Use the Lambda function to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Email Service (Amazon SES).
- C. Schedule an AWS Lambda function to start a manual image scan every hour. Configure Amazon EventBridge to invoke another Lambda function when a scan is complete. Use the second Lambda function to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
- D. Configure periodic image scan on the repository. Configure scan results to be added to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Step Functions state machine when a new message is added to the SQS queue. Use the Step Functions state machine to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Email Service (Amazon SES).

Answer: A

Explanation: <https://docs.aws.amazon.com/AmazonECR/latest/userguide/ecr-eventbridge.html> "Activating an

AWS Step Functions state machine" <https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-creating-lambda-state-machine.html>

Question: 235

A company wants to use AWS for disaster recovery for an on-premises application. The company has hundreds

of Windows-based servers that run the application. All the servers mount a common share. The company has an RTO of 15 minutes and an RPO of 5 minutes. The solution must support native failover and fallback capabilities.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an AWS Storage Gateway File Gateway. Schedule daily Windows server backups. Save the data to Amazon S3. During a disaster, recover the on-premises servers from the backup. During failback, run the on-premises servers on Amazon EC2 instances.
- B. Create a set of AWS CloudFormation templates to create infrastructure. Replicate all data to Amazon Elastic File System (Amazon EFS) by using AWS DataSync. During a disaster, use AWS CodePipeline to deploy the templates to restore the on-premises servers. Fail back the data by using DataSync.
- C. Create an AWS Cloud Development Kit (AWS CDK) pipeline to stand up a multi-site active-active environment on AWS. Replicate data into Amazon S3 by using the s3 sync command. During a disaster, swap DNS endpoints to point to AWS. Fail back the data by using the s3 sync command.
- D. Use AWS Elastic Disaster Recovery to replicate the on-premises servers. Replicate data to an Amazon FSx for Windows File Server file system by using AWS DataSync. Mount the file system to AWS servers. During a disaster, fail over the on-premises servers to AWS. Fail back to new or existing servers by using Elastic Disaster Recovery.

Answer: D

Explanation:

Question: 236

A solutions architect needs to define a reference architecture for a solution for three-tier applications with web application, and NoSQL data layers. The reference architecture must meet the following requirements:

- High availability within an AWS Region
- Able to fail over in 1 minute to another AWS Region for disaster recovery
- Provide the most efficient solution while minimizing the impact on the user experience

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

- A. Use an Amazon Route 53 weighted routing policy set to 100/0 across the two selected Regions. Set Time to Live (TTL) to 1 hour.
- B. Use an Amazon Route 53 failover routing policy for failover from the primary Region to the disaster recovery Region. Set Time to Live (TTL) to 30 seconds.
- C. Use a global table within Amazon DynamoDB so data can be accessed in the two selected Regions.
- D. Back up data from an Amazon DynamoDB table in the primary Region every 60 minutes and then write the data to Amazon S3. Use S3 Cross-Region replication to copy the data from the primary Region to the disaster recovery Region. Have a script import the data into DynamoDB in a disaster recovery scenario.
- E. Implement a hot standby model using Auto Scaling groups for the web and application layers across multiple Availability Zones in the Regions. Use zonal Reserved Instances for the minimum number of servers and On-Demand Instances for any additional resources.
- F. Use Auto Scaling groups for the web and application layers across multiple Availability Zones in the Regions. Use Spot Instances for the required resources.

Answer: B,C,E

Explanation:

The requirements can be achieved by using an Amazon DynamoDB database with a global table. DynamoDB is a NoSQL database so it fits the requirements. A global table also allows both reads and writes to occur in both Regions. For the web and application tiers Auto Scaling groups should be configured. Due to the 1-minute RTO these must be configured in an active/passive state. The best pricing model to lower price but ensure resources are available when needed is to use a combination of zonal reserved instances and on-demand instances. To failover between the Regions, a Route 53 failover routing policy can be configured with a TTL configured on the record of 30 seconds. This will mean clients must resolve against Route 53 every 30 seconds to get the latest record. In a failover scenario the clients would be redirected to the secondary site if the primary site is unhealthy.

Question: 237

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

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

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

Answer: B,E,F

Explanation:

"This operation can be called only from the organization's management account. Member accounts can remove themselves with LeaveOrganization instead." https://docs.aws.amazon.com/organizations/latest/APIReference/API_RemoveAccountFromOrganization.html

Question: 238

A company wants to containerize a multi-tier web application and move the application from an on-premises data center to AWS. The application includes web, application, and database tiers. The company needs to make the application fault tolerant and scalable. Some frequently accessed data must always be available across application servers. Frontend web servers need session persistence and must scale to meet increases in traffic.

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

- A. Run the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. Use Amazon Elastic File System (Amazon EFS) for data that is frequently accessed between the web and application tiers. Store the frontend web server session data in Amazon Simple Queue Service (Amazon SQS).
- B. Run the application on Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use Amazon ElastiCache for Redis to cache frontend web server session data. Use Amazon Elastic Block Store (Amazon EBS) with Multi-Attach on EC2 instances that are distributed across multiple Availability Zones.
- C. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node groups. Use ReplicaSets to run the web servers and applications. Create an Amazon Elastic File System (Amazon EFS) Me system. Mount the EFS file system across all EKS pods to store frontend web server session data.
- D. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) Configure Amazon EKS to use managed node groups. Run the web servers and application as Kubernetes deployments in the EKS cluster. Store the frontend web server session data in an Amazon DynamoDB table. Create an Amazon Elastic File System (Amazon EFS) volume that all applications will mount at the time of deployment.

Answer: D

Explanation:

Deploying the application on Amazon EKS with managed node groups simplifies the operational overhead of managing the Kubernetes cluster. Running the web servers and application as Kubernetes deployments ensures that the desired number of pods are always running and can scale up or down as needed. Storing the frontend web server session data in an Amazon DynamoDB table provides a fast, scalable, and durable storage option that can be accessed across multiple Availability Zones. Creating an Amazon EFS volume that all applications will mount at the time of deployment allows the application to share data that is frequently accessed between the web and application tiers. Reference: <https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html><https://docs.aws.amazon.com/eks/latest/userguide/deployments.html>

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html><https://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html>

Question: 239

A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endmost IP addresses are not permitted.

The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.

Which solution meets these requirements?

- A. Register the customer-owned block of IP addresses in the company's AWS account. Create Elastic IP addresses from the address pool and assign them to an AWS Transfer for SFTP endpoint. Use AWS Transfer to store the files in Amazon S3.
- B. Add a subnet containing the customer-owned block of IP addresses to a VPC. Create Elastic IP addresses from the address pool and assign them to an Application Load Balancer (ALB). Launch EC2 instances hosting FTP services in an Auto Scaling group behind the ALB. Store the files in attached Amazon Elastic Block Store (Amazon EBS) volumes.
- C. Register the customer-owned block of IP addresses with Amazon Route 53. Create alias records in Route 53 that point to a Network Load Balancer (NLB). Launch EC2 instances hosting FTP services in an Auto Scaling group behind the NLB. Store the files in Amazon S3.
- D. Register the customer-owned block of IP addresses in the company's AWS account. Create Elastic IP addresses from the address pool and assign them to an Amazon S3 VPC endpoint. Enable SFTP support on the S3 bucket.

Answer: A

Explanation:

Bring your own IP addresses (BYOIP) You can bring part or all of your publicly routable IPv4 or IPv6 address range from your on-premises network to your AWS account. You continue to own the address range, but AWS advertises it on the internet by default. After you bring the address range to AWS, it appears in your AWS account as an address pool.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-byoip.html> AWS Transfer for SFTP enables you to easily move your file transfer workloads that use the Secure Shell File Transfer Protocol (SFTP) to AWS without needing to modify your applications or manage any SFTP servers. <https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/>

Question: 240

A company is designing an AWS Organizations structure. The company wants to standardize a process to apply tags across the entire organization. The company will require tags with specific values when a user creates a new resource. Each of the company's OUs will have unique tag values.

Which solution will meet these requirements?

- A. Use an SCP to deny the creation of resources that do not have the required tags. Create a tag policy that includes the tag values that the company has assigned to each OU. Attach the tag policies to the OUs.
- B. Use an SCP to deny the creation of resources that do not have the required tags. Create a tag policy that includes the tag values that the company has assigned to each OU. Attach the tag policies to the organization's management account.
- C. Use an SCP to allow the creation of resources only when the resources have the required tags. Create a tag policy that includes the tag values that the company has assigned to each OU. Attach the tag policies to the

OUS.

- D. Use an SCP to deny the creation of resources that do not have the required tags. Define the list of tags. Attach the SCP to the OUs

Answer: A

Explanation:

<https://aws.amazon.com/blogs/mt/implement-aws-resource-tagging-strategy-using-aws-tag-policies-and-service-control-policies-scps/>

Question: 241

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

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

Answer: B

Explanation:

AWS DMS can migrate data from a source database to a target database in AWS, using change data capture (CDC) to replicate ongoing changes and keep the databases in sync. Setting Amazon S3 as a target allows storing the migrated data in a durable and cost-effective storage service. When the source and destination are fully synchronized, the data can be loaded from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB instance, which is a managed database service that simplifies database administration tasks.

Reference:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.SQLServer.htmlhttps://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServer.html

Question: 242

A company has set up its entire infrastructure on AWS. The company uses Amazon EC2 instances to host its ecommerce website and uses Amazon S3 to store static data.

a. Three engineers at the company handle the cloud administration and development through one AWS account. Occasionally, an engineer alters an EC2 security group configuration of another engineer and causes noncompliance issues in the environment.

A solutions architect must set up a system that tracks changes that the engineers make. The system must send alerts when the engineers make noncompliant changes to the security settings for the EC2 instances. What is the FASTEST way for the solutions architect to meet these requirements?

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

Answer: D

Explanation:

<https://aws.amazon.com/es/blogs/industries/how-to-monitor-alert-and-remediate-non-compliant-hipaa-findings-on-aws/>

Question: 243

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? (Select 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 Cloud Formation template, write a condition that updates the network routes when a replacement instance is launched.

Answer: B,D,E

Explanation:

Question: 244

A company is migrating its development and production workloads to a new organization in AWS Organizations. The company has created a separate member account for development and a separate member account for production. Consolidated billing is linked to the management account. In the management account, a solutions architect needs to create an IAM user that can stop or terminate resources in both member accounts. Which solution will meet this requirement?

- A. Create an IAM user and a cross-account role in the management account. Configure the crossaccount role with least privilege access to the member accounts.
- B. Create an IAM user in each member account. In the management account, create a cross-account role that has least privilege access. Grant the IAM users access to the cross-account role by using a trust policy.
- C. Create an IAM user in the management account. In the member accounts, create an IAM group that has least privilege access. Add the IAM user from the management account to each IAM group in the member accounts.
- D. Create an IAM user in the management account. In the member accounts, create cross-account roles that have least privilege access. Grant the IAM user access to the roles by using a trust policy.

Answer: D

Explanation:

Cross account role should be created in destination(member) account. The role has trust entity to master account.

Question: 245

A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.

After initial deployment, the company observes 1.000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.

Which approach should the company take to secure its API?

- A. Create an Amazon CloudFront distribution with the API as the origin. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per day. Associate the web ACL with the CloudFront distribution. Configure CloudFront with an origin access identity (OAI) and associate it with the distribution. Configure API Gateway to ensure only the OAI can run the POST method.
- B. Create an Amazon CloudFront distribution with the API as the origin. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per day. Associate the web ACL with the CloudFront distribution. Add a custom header to the CloudFront distribution populated with an API key. Configure the API to require an API key on the POST method.
- C. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners. Associate the web ACL with the API. Create a resource policy with a request limit and associate it with the API. Configure the API to require an API key on the POST method.
- D. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners. Associate the web ACL with the API. Create a usage plan with a request limit and associate it with the API. Create an API key and add it to the usage plan.

Answer: D

Explanation:

"A usage plan specifies who can access one or more deployed API stages and methods—and also how much and how fast they can access them. The plan uses API keys to identify API clients and meters access to the associated API stages for each key. It also lets you configure throttling limits and quota limits that are enforced on individual client API keys." <https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html>

A rate-based rule tracks the rate of requests for each originating IP address, and triggers the rule action on IPs with rates that go over a limit. You set the limit as the number of requests per 5-minute time span. The following caveats apply to AWS WAF rate-based rules: The minimum rate that you can set is 100. AWS WAF checks the rate of requests every 30 seconds, and counts requests for the prior five minutes each time. Because of this, it's possible for an IP address to send requests at too high a rate for 30 seconds before AWS WAF detects and blocks it. AWS WAF can block up to 10,000 IP addresses. If more than 10,000 IP addresses send high rates of requests at the same time, AWS WAF will only block 10,000 of them.

"<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html>

Question: 246

A company provides auction services for artwork and has users across North America and Europe. The company hosts its application in Amazon EC2 instances in the us-east-1 Region. Artists upload photos of their work as large-size, high-resolution image files from their mobile phones to a centralized Amazon S3 bucket created in the us-east-1 Region. The users in Europe are reporting slow performance for their Image uploads.

How can a solutions architect improve the performance of the image upload process?

- A. Redeploy the application to use S3 multipart uploads.
- B. Create an Amazon CloudFront distribution and point to the application as a custom origin

- C. Configure the buckets to use S3 Transfer Acceleration.
- D. Create an Auto Scaling group for the EC2 instances and create a scaling policy.

Answer: C

Explanation:

Transfer acceleration. S3 Transfer Acceleration utilizes the Amazon CloudFront global network of edge locations to accelerate the transfer of data to and from S3 buckets. By enabling S3 Transfer Acceleration on the centralized S3 bucket, the users in Europe will experience faster uploads as their data will be routed through the closest CloudFront edge location.

Question: 247

A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.

When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.

What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

- A. Configure Amazon EventBridge to invoke an AWS Lambda function when a new EC2 instance is launched into the Auto Scaling group. Code the Lambda function to associate the EC2 instances with the CodeDeploy deployment group.
- B. Write a script to suspend Amazon EC2 Auto Scaling operations before the deployment of new code. When the deployment is complete, create a new AMI and configure the Auto Scaling group's launch template to use the new AMI for new launches. Resume Amazon EC2 Auto Scaling operations.
- C. Create a new AWS CodeBuild project that creates a new AMI that contains the new code. Configure CodeBuild to update the Auto Scaling group's launch template to the new AMI. Run an Amazon EC2 Auto Scaling instance refresh operation.
- D. Create a new AMI that has the CodeDeploy agent installed. Configure the Auto Scaling group's launch template to use the new AMI. Associate the CodeDeploy deployment group with the Auto Scaling group instead of the EC2 instances.

Answer: D

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html>

Question: 248 A company is migrating a legacy application from an on-premises data center to AWS. The application uses MongoDB as a key-value database. According to the company's technical guidelines, all Amazon EC2 instances must be hosted in a private subnet without an internet connection. In addition, all connectivity between applications and databases must be encrypted. The database must be able to scale based on demand.

Which solution will meet these requirements?

- A. Create new Amazon DocumentDB (with MongoDB compatibility) tables for the application with Provisioned IOPS volumes. Use the instance endpoint to connect to Amazon DocumentDB.
- B. Create new Amazon DynamoDB tables for the application with on-demand capacity. Use a gateway VPC endpoint for DynamoDB to connect to the DynamoDB tables.
- C. Create new Amazon DynamoDB tables for the application with on-demand capacity. Use an interface VPC endpoint for DynamoDB to connect to the DynamoDB tables.
- D. Create new Amazon DocumentDB (with MongoDB compatibility) tables for the application with Provisioned IOPS volumes. Use the cluster endpoint to connect to Amazon DocumentDB.

Answer: A

Explanation:

A is the correct answer because it uses Amazon DocumentDB (with MongoDB compatibility) as a key-value database that can scale based on demand and supports encryption in transit and at rest. Amazon DocumentDB is a fully managed document database service that is designed to be compatible with the MongoDB API. It is a NoSQL database that is optimized for storing, indexing, and querying JSON data. Amazon DocumentDB supports encryption in transit using TLS and encryption at rest using AWS Key Management Service (AWS KMS). Amazon DocumentDB also supports provisioned IOPS volumes that can scale up to 64 TiB of storage and 256,000 IOPS per cluster. To connect to Amazon DocumentDB, you can use the instance endpoint, which connects to a specific instance in the cluster, or the cluster endpoint, which connects to the primary instance or one of the replicas in the cluster. Using the cluster endpoint is recommended for high availability and load balancing purposes. Reference: <https://docs.aws.amazon.com/documentdb/latest/developerguide/what-is.html><https://docs.aws.amazon.com/documentdb/latest/developerguide/security.encryption.html><https://docs.aws.amazon.com/documentdb/latest/developerguide/limits.html><https://docs.aws.amazon.com/documentdb/latest/developerguide/connecting.html>

Question: 249

A company is developing a new on-demand video application that is based on microservices. The application will have 5 million users at launch and will have 30 million users after 6 months. The company has deployed the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. The company developed the application by using ECS services that use the HTTPS protocol. A solutions architect needs to implement updates to the application by using blue/green deployments. The solution must distribute traffic to each ECS service through a load balancer. The application must automatically adjust the number of tasks in response to an Amazon CloudWatch alarm.

Which solution will meet these requirements?

- www.atmicnetworks.com
- A. Configure the ECS services to use the blue/green deployment type and a Network Load Balancer. Request increases to the service quota for tasks per service to meet the demand.
 - B. Configure the ECS services to use the blue/green deployment type and a Network Load Balancer. Implement an Auto Scaling group for each ECS service by using the Cluster Autoscaler.
 - C. Configure the ECS services to use the blue/green deployment type and an Application Load Balancer. Implement an Auto Seating group for each ECS service by using the Cluster Autoscaler.
 - D. Configure the ECS services to use the blue/green deployment type and an Application Load Balancer. Implement Service Auto Scaling for each ECS service.

Answer: C

Explanation:

<https://repost.aws/knowledge-center/ecs-fargate-service-auto-scaling>

Question: 250

A company wants to migrate to AWS. The company is running thousands of VMs in a VMware ESXi environment. The company has no configuration management database and has little Knowledge about the utilization of the VMware portfolio.

A solutions architect must provide the company with an accurate inventory so that the company can plan for a cost-effective migration.

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

- A. Use AWS Systems Manager Patch Manager to deploy Migration Evaluator to each VM. Review the collected data in Amazon QuickSight. Identify servers that have high utilization. Remove the servers that have high utilization from the migration list. Import the data to AWS Migration Hub.
- B. Export the VMware portfolio to a csv file. Check the disk utilization for each server. Remove servers that have high utilization. Export the data to AWS Application Migration Service. Use AWS Server Migration Service (AWS SMS) to migrate the remaining servers.
- C. Deploy the Migration Evaluator agentless collector to the ESXi hypervisor. Review the collected data in Migration Evaluator. Identify inactive servers. Remove the inactive servers from the migration list. Import the data to AWS Migration Hub.
- D. Deploy the AWS Application Migration Service Agent to each VM. When the data is collected, use Amazon Redshift to import and analyze the data. Use Amazon QuickSight for data visualization.

Answer: C

Explanation:

<https://aws.amazon.com/migration-evaluator/features/>

Question: 251

A solutions architect at a large company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the internet. The company deploys resources only into a single AWS Region.

The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.

Which solution meets these requirements?

- A. Create a new VPC for outbound traffic to the internet. Connect the existing transit gateway to the new VPC. Configure a new NAT gateway. Create an Auto Scaling group of Amazon EC2 instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Region. **Modify all default routes to point to the proxy's Auto Scaling group.**
- B. Create a new VPC for outbound traffic to the internet. Connect the existing transit gateway to the new VPC. Configure a new NAT gateway. Use an AWS Network Firewall firewall for rule-based filtering. Create Network Firewall endpoints in each Availability Zone. **Modify all default routes to point to the Network Firewall endpoints.**
- C. Create an AWS Network Firewall firewall for rule-based filtering in each AWS account. **Modify all default routes to point to the Network Firewall firewalls in each account.**
- D. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering. **Modify all default routes to point to the proxy's Auto Scaling group.**

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/deployment-models-for-aws-network-firewall/>

Question: 252

A company has an on-premises Microsoft SQL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS.

Which solution meets these requirements MOST cost-effectively?

- A. Create a new S3 bucket. Deploy an AWS Storage Gateway file gateway within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to the new SMB file share.
- B. Create an Amazon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the

- Direct Connect connection. Create a new SMB file share. Write nightly database exports to an SMB file share on the Amazon FSx file system. Enable nightly backups.
- C. Create an Amazon FSx for Windows File Server Multi-AZ file system within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to an SMB file share on the Amazon FSx file system. Enable nightly backups.
- D. Create a new S3 bucket. Deploy an AWS Storage Gateway volume gateway within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.

Answer: A

Explanation:

<https://docs.aws.amazon.com/filegateway/latest/files3/CreatingAnSMBFileShare.html>

Question: 253

A company is running an application on Amazon EC2 instances in the AWS Cloud. The application is using a MongoDB database with a replica set as its data tier. The MongoDB database is installed on systems in the company's on-premises data center and is accessible through an AWS Direct Connect connection to the data center environment.

A solutions architect must migrate the on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility).

Which strategy should the solutions architect choose to perform this migration?

- A. Create a fleet of EC2 instances. Install MongoDB Community Edition on the EC2 instances, and create a database. Configure continuous synchronous replication with the database that is running in the on-premises data center.
- B. Create an AWS Database Migration Service (AWS DMS) replication instance. Create a source endpoint for the on-premises MongoDB database by using change data capture (CDC). Create a target endpoint for the Amazon DocumentDB database. Create and run a DMS migration task.
- C. Create a data migration pipeline by using AWS Data Pipeline. Define data nodes for the on-premises MongoDB database and the Amazon DocumentDB database. Create a scheduled task to run the data pipeline.
- D. Create a source endpoint for the on-premises MongoDB database by using AWS Glue crawlers. Configure continuous asynchronous replication between the MongoDB database and the Amazon DocumentDB database.

Answer: B

Explanation:

<https://aws.amazon.com/getting-started/hands-on/move-to-managed/migrate-mongodb-to-documentdb/>

Question: 254

A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts. Each VPC consists of 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:

https://d1.awsstatic.com/architecture-diagrams/ArchitectureDiagrams/NAT-gateway-centralized-egress-ra.pdf?did=wp_card&trk=wp_card

Question: 255

A company needs to migrate its customer transactions database from on premises to AWS. The database resides on an Oracle DB instance that runs on a Linux server. According to a new security requirement, the company must rotate the database password each year.

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

- A. Convert the database to Amazon DynamoDB by using the AWS Schema Conversion Tool (AWS SCT). Store the password in AWS Systems Manager Parameter Store. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.
- B. Migrate the database to Amazon RDS for Oracle. Store the password in AWS Secrets Manager. Turn ON automatic rotation. Configure a yearly rotation schedule.
- C. Migrate the database to an Amazon EC2 instance. Use AWS Systems Manager Parameter Store to keep and rotate the connection string by using an AWS Lambda function on a yearly schedule.
- D. Migrate the database to Amazon Neptune by using the AWS Schema Conversion Tool (AWS SCT). Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.

Answer: B

Explanation:

Question: 256

A financial services company loaded millions of historical stock trades into an Amazon DynamoDB table. The table uses on-demand capacity mode. Once each day at midnight, a few million new records are loaded into the table. Application read activity against the table happens in bursts throughout the day, and a limited set of keys are repeatedly looked up. The company needs to reduce costs associated with DynamoDB.

Which strategy should a solutions architect recommend to meet this requirement?

- A. Deploy an Amazon ElastiCache cluster in front of the DynamoDB table.
- B. Deploy DynamoDB Accelerator (DAX). Configure DynamoDB auto scaling. Purchase Savings Plans in Cost Explorer
- C. Use provisioned capacity mode. Purchase Savings Plans in Cost Explorer.
- D. Deploy DynamoDB Accelerator (DAX). Use provisioned capacity mode. Configure DynamoDB auto scaling.

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html#HowItWorks.ProvisionedThroughput.Manual>

Question: 257

A company wants to manage the costs associated with a group of 20 applications that are infrequently used, but are still business-critical, by migrating to AWS. The applications are a mix of Java and Node.js spread across different instance clusters. The company wants to minimize costs while standardizing by using a single deployment methodology.

Most of the applications are part of month-end processing routines with a small number of concurrent users, but they are occasionally run at other times. Average application memory consumption is less than 1 GB, though some applications use as much as 2.5 GB of memory during peak processing. The most important application in the group is a billing report written in Java that accesses multiple data sources and often runs for several hours.

Which is the MOST cost-effective solution?

- A. Deploy a separate AWS Lambda function for each application. Use AWS CloudTrail logs and Amazon CloudWatch alarms to verify completion of critical jobs.
- B. Deploy Amazon ECS containers on Amazon EC2 with Auto Scaling configured for memory utilization of 75%. Deploy an ECS task for each application being migrated with ECS task scaling. Monitor services and hosts by using Amazon CloudWatch.

- C. Deploy AWS Elastic Beanstalk for each application with Auto Scaling to ensure that all requests have sufficient resources. Monitor each AWS Elastic Beanstalk deployment by using CloudWatch alarms.
- D. Deploy a new Amazon EC2 instance cluster that co-hosts all applications by using EC2 Auto Scaling and Application Load Balancers. Scale cluster size based on a custom metric set on instance memory utilization. Purchase 3-year Reserved Instance reservations equal to the GroupMaxSize parameter of the Auto Scaling group.

Answer: B

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.cloudwatch.html> Elastic Beanstalk automatically uses Amazon CloudWatch to help you monitor your application and environment status. You can navigate to the Amazon CloudWatch console to see your dashboard and get an overview of all of your resources as well as your alarms. You can also choose to view more metrics or add custom metrics.

Question: 258

During an audit, a security team discovered that a development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository. The security team wants to automatically find and remediate instances of this security vulnerability.

Which solution will ensure that the credentials are appropriately secured automatically?

- A. Run a script nightly using AWS Systems Manager Run Command to search for credentials on the development instances. If found, use AWS Secrets Manager to rotate the credentials.
- B. Use a scheduled AWS Lambda function to download and scan the application code from CodeCommit. If credentials are found, generate new credentials and store them in AWS KMS.
- C. Configure Amazon Macie to scan for credentials in CodeCommit repositories. If credentials are found, trigger an AWS Lambda function to disable the credentials and notify the user.
- D. Configure a CodeCommit trigger to invoke an AWS Lambda function to scan new code submissions for credentials. If credentials are found, disable them in AWS IAM and notify the user.

Answer: D

Explanation:

CodeCommit may use S3 on the back end (and it also uses DynamoDB on the back end) but I don't think they're stored in buckets that you can see or point Macie to. In fact, there are even solutions out there describing how to copy your repo from CodeCommit into S3 to back it up:
<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/automate-event-driven-backups-from-codecommit-to-amazon-s3-using-codebuild-and-cloudwatch-events.html>

Question: 259

A company is rearchitecting its applications to run on AWS. The company's infrastructure includes multiple Amazon EC2 instances. The company's development team needs different levels of access. The company wants to implement a policy that requires all Windows EC2 instances to be joined to an Active Directory domain on AWS. The company also wants to implement enhanced security processes such as multi-factor authentication (MFA). The company wants to use managed AWS services wherever possible.

Which solution will meet these requirements?

- A. Create an AWS Directory Service for Microsoft Active Directory implementation. Launch an Amazon Workspace. Connect to and use the Workspace for domain security configuration tasks.
- B. Create an AWS Directory Service for Microsoft Active Directory implementation. Launch an EC2 instance. Connect to and use the EC2 instance for domain security configuration tasks.
- C. Create an AWS Directory Service Simple AD implementation. Launch an EC2 instance. Connect to and use the EC2 instance for domain security configuration tasks.
- D. Create an AWS Directory Service Simple AD implementation. Launch an Amazon Workspace. Connect to and use the Workspace for domain security configuration tasks.

Answer: A

Explanation:

A is the correct answer because it uses AWS Directory Service for Microsoft Active Directory to join the Windows EC2 instances to an Active Directory domain on AWS and enable MFA. AWS Directory Service for Microsoft Active Directory, also known as AWS Managed Microsoft AD, is a fully managed service that is powered by Windows Server 2019. It allows you to run directory-aware workloads in the AWS Cloud, including Microsoft SharePoint and custom .NET and SQL Server-based applications. You can also configure a trust relationship between AWS Managed Microsoft AD in the AWS Cloud and your existing on-premises Microsoft Active Directory. AWS Managed Microsoft AD supports MFA by integrating with your existing RADIUS-based MFA infrastructure. To join the Windows EC2 instances to an Active Directory domain on AWS, you can use an Amazon Workspace, which is a fully managed, secure desktop computing service that runs on AWS. You can connect to and use the Workspace for domain security configuration tasks. Reference:

https://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_microsoft_ad.html<https://docs.aws.amazon.com/workspaces/latest/adminguide/amazon-workspaces.html>

Question: 260

A company is running an application in the AWS Cloud. The application consists of microservices that run on a fleet of Amazon EC2 instances in multiple Availability Zones behind an Application Load Balancer. The company recently added a new REST API that was implemented in Amazon API Gateway. Some of the older microservices that run on EC2 instances need to call this new API. The company does not want the API to be accessible from the public internet and does not want proprietary data to traverse the public internet.

What should a solutions architect do to meet these requirements?

- A. Create an AWS Site-to-Site VPN connection between the VPC and the API Gateway. Use API Gateway to generate a unique API key for each microservice. Configure the API methods to require the key.
- B. Create an interface VPC endpoint for API Gateway, and set an endpoint policy to only allow access to the specific API. Add a resource policy to API Gateway to only allow access from the VPC endpoint. Change the API Gateway endpoint type to private.
- C. Modify the API Gateway to use IAM authentication. Update the IAM policy for the IAM role that is assigned to the EC2 Instances to allow access to the API Gateway. Move the API Gateway into a new VPC. Deploy a transit gateway and connect the VPCs.
- D. Create an accelerator in AWS Global Accelerator, and connect the accelerator to the API Gateway. Update the route table for all VPC subnets with a route to the created Global Accelerator endpoint IP address. Add an API key for each service to use for authentication.

Answer: B

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-vpc-endpoint-policies.html>

Question: 261

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:

Because MSK has a maximum number of client connections of 1000 per second and the company has 10,000 sensors, the MSK likely will not be able to handle all

connectionshttps://docs.aws.amazon.com/msk/latest/developerguide/limits.html

Question: 262

A company is developing a gene reporting device that will collect genomic information to assist researchers with collecting large samples of data from a diverse population. The device will push 8 KB of genomic data every second to a data platform that will need to process and analyze the data and provide information back to researchers. The data platform must meet the following requirements:

- Provide near-real-time analytics of the inbound genomic data
 - Ensure the data is flexible, parallel, and durable
 - Deliver results of processing to a data warehouse
- Which strategy should a solutions architect use to meet these requirements?

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

Answer: B

Explanation:

Kinesis Data Streams is a real-time streaming service and provide near-real-time analytics. Also the question "Deliver results of processing to a data warehouse" and this option has redshift cluster which is a powerful data warehousing solution that can handle large-scale analytics workloads.

Question: 263

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? (Select 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:

B:<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html#HowItWorks.requests>

D:<https://aws.amazon.com/blogs/compute/robust-serverless-application-design-with-aws-lambda-dlq/c>

Question: 264

A company runs many workloads on AWS and uses AWS Organizations to manage its accounts. The workloads are hosted on Amazon EC2, AWS Fargate, and AWS Lambda.

a. Some of the workloads have unpredictable demand. Accounts record high usage in some months and low usage in other months.

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

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

A. Purchase Reserved Instances for the organization to match the size and number of the most common EC2 instances from the member accounts.

B. Purchase a Compute Savings Plan for the organization from the management account by using the recommendation at the management account level.

C. Purchase Reserved Instances for each member account that had high EC2 usage according to the data from the last 6 months.

D. Purchase an EC2 Instance Savings Plan for each member account from the management account based on EC2 usage data from the last 6 months.

Answer: B

Explanation:

Question: 265

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

a. The company needs to decouple the product data from the user session data. The company also needs to implement replication in another AWS Region for disaster recovery.

Which solution will meet these requirements with the HIGHEST performance?

A. Create an Amazon RDS DB instance with separate schemas to host the product data and the user session data. Configure a read replica for the DB instance in another Region.

B. Create an Amazon RDS DB instance to host the product data. Configure a read replica for the DB instance in another Region. Create a global datastore in Amazon ElastiCache for Memcached to host the user session data.

C. Create two Amazon DynamoDB global tables. Use one global table to host the product data. Use the other global table to host the user session data. Use DynamoDB Accelerator (DAX) for caching.

D. Create an Amazon RDS DB instance to host the product data. Configure a read replica for the DB instance in another Region. Create an Amazon DynamoDB global table to host the user session data

Answer: B

Explanation:

Question: 266

A company uses AWS Organizations to manage its AWS accounts. The company needs a list of all its Amazon EC2 instances that have underutilized CPU or memory usage. The company also needs recommendations for how to downsize these underutilized instances.

Which solution will meet these requirements with the LEAST effort?

- A. Install a CPU and memory monitoring tool from AWS Marketplace on all the EC2 Instances. Store the findings in Amazon S3. Implement a Python script to identify underutilized instances. Reference EC2 instance pricing information for recommendations about downsizing options.
- B. Install the Amazon CloudWatch agent on all the EC2 instances by using AWS Systems Manager. Retrieve the resource optimization recommendations from AWS Cost Explorer in the organization's management account. Use the recommendations to downsize underutilized instances in all accounts of the organization.
- C. Install the Amazon CloudWatch agent on all the EC2 instances by using AWS Systems Manager. Retrieve the resource optimization recommendations from AWS Cost Explorer in each account of the organization. Use the recommendations to downsize underutilized instances in all accounts of the organization.
- D. Install the Amazon CloudWatch agent on all the EC2 instances by using AWS Systems Manager. Create an AWS Lambda function to extract CPU and memory usage from all the EC2 instances. Store the findings as files in Amazon S3. Use Amazon Athena to find underutilized instances. Reference EC2 instance pricing information for recommendations about downsizing options.

Answer: B

Explanation:

Question: 267

A company needs to create and manage multiple AWS accounts for a number of departments from a central location. The security team requires read-only access to all accounts from its own AWS account. The company is using AWS Organizations and created an account for the security team. How should a solutions architect meet these requirements?

- A. Use the OrganizationAccountAccessRole IAM role to create a new IAM policy with read-only access in each member account. Establish a trust relationship between the IAM policy in each member account and the security account. Ask the security team to use the IAM policy to gain access.
- B. Use the Organization AccountAccessRole IAM role to create a new IAM role with read-only access in each member account. Establish a trust relationship between the IAM role in each member account and the security account. Ask the security team to use the IAM role to gain access.

- C. Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the OrganizationAccountAccessRole IAM role in the management account from the security account. Use the generated temporary credentials to gain access.
- D. Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the OrganizationAccountAccessRole IAM role in the member account from the security account. Use the generated temporary credentials to gain access.

Answer: B

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html#orgs_manage_accounts_access-cross-account-role "When you create a member account using the AWS Organizations console, AWS Organizations automatically creates an IAM role named OrganizationAccountAccessRole in the account" you need OrganizationAccountAccessRole in member account to create an read-only role and use role from security team to assume this readonly role.

Question: 268

A company runs an unauthenticated static website (www.example.com) that includes a registration form for users. The website uses Amazon S3 for hosting and uses Amazon CloudFront as the content delivery network with AWS WAF configured. When the registration form is submitted, the website calls an Amazon API Gateway API endpoint that invokes an AWS Lambda function to process the payload and forward the payload to an external API call.

During testing, a solutions architect encounters a cross-origin resource sharing (CORS) error. The solutions architect confirms that the CloudFront distribution origin has the Access-Control-Allow-Origin header set to www.example.com.

What should the solutions architect do to resolve the error?

- A. Change the CORS configuration on the S3 bucket. Add rules for CORS to the Allowed Origin element for www.example.com.
- B. Enable the CORS setting in AWS WAF. Create a web ACL rule in which the Access-Control-Allow-Origin header is set to www.example.com.
- C. Enable the CORS setting on the API Gateway API endpoint. Ensure that the API endpoint is configured to return all responses that have the Access-Control-Allow-Origin header set to www.example.com.
- D. Enable the CORS setting on the Lambda function. Ensure that the return code of the function has the Access-Control-Allow-Origin header set to www.example.com.

Answer: C

Explanation:

CORS errors occur when a web page hosted on one domain tries to make a request to a server hosted on another domain. In this scenario, the registration form hosted on the static website is trying to make a request to the API Gateway API endpoint hosted on a different domain, which is causing the error. To resolve this error,

the Access-Control-Allow-Origin header needs to be set to the domain from which the request is being made. In this case, the header is already set to www.example.com on the CloudFront distribution origin. Therefore, the solutions architect should enable the CORS setting on the API Gateway API endpoint and ensure that the API endpoint is configured to return all responses that have the Access-Control-Allow-Origin header set to www.example.com. This will allow the API endpoint to respond to requests from the static website without a CORS error. <https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-cors-errors/>

Question: 269

A large education company recently introduced Amazon Workspaces to provide access to internal applications across multiple universities. The company is storing user profiles on an Amazon FSx (or Windows File Server file system). The file system is configured with a DNS alias and is connected to a self-managed Active Directory. As more users begin to use the Workspaces, login time increases to unacceptable levels.

An investigation reveals a degradation in performance of the file system. The company created the file system on HDD storage with a throughput of 16 MBps. A solutions architect must improve the performance of the file system during a defined maintenance window.

What should the solutions architect do to meet these requirements with the LEAST administrative effort?

- A. Use AWS Backup to create a point-in-time backup of the file system. Restore the backup to a new FSx for Windows File Server file system. Select SSD as the storage type. Select 32 MBps as the throughput capacity. When the backup and restore process is completed, adjust the DNS alias accordingly. Delete the original file system.
- B. Disconnect users from the file system. In the Amazon FSx console, update the throughput capacity to 32 MBps. Update the storage type to SSD. Reconnect users to the file system.
- C. Deploy an AWS DataSync agent onto a new Amazon EC2 Instance. Create a task. Configure the existing file system as the source location. Configure a new FSx for Windows File Server file system with SSD storage and 32 MBps of throughput as the target location. Schedule the task. When the task is completed, adjust the DNS alias accordingly. Delete the original file system.
- D. Enable shadow copies on the existing file system by using a Windows PowerShell command. Schedule the shadow copy job to create a point-in-time backup of the file system. Choose to restore previous versions. Create a new FSx for Windows File Server file system with SSD storage and 32 MBps of throughput. When the copy job is completed, adjust the DNS alias. Delete the original file system.

Answer: C

Explanation:

<https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html>

Question: 270

A company has an application that is deployed on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are part of an Auto Scaling group. The application has unpredictable workloads and

frequently scales out and in. The company's development team wants to analyze application logs to find ways to improve the application's performance. However, the logs are no longer available after instances scale in.

Which solution will give the development team the ability to view the application logs after a scale in event?

- A. Enable access logs for the ALB. Store the logs in an Amazon S3 bucket.
- B. Configure the EC2 instances to publish logs to Amazon CloudWatch Logs by using the unified CloudWatch agent.
- C. Modify the Auto Scaling group to use a step scaling policy.
- D. Instrument the application with AWS X-Ray tracing.

Answer: B

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

Question: 271

A company has an application that generates reports and stores them in an Amazon S3 bucket. When a user accesses their report, the application generates a signed URL to allow the user to download the report. The company's security team has discovered that the files are public and that anyone can download them without authentication. The company has suspended the generation of new reports until the problem is resolved. Which set of actions will immediately remediate the security issue without impacting the application's normal workflow?

- A. Create an AWS Lambda function that applies a deny all policy for users who are not authenticated. Create a scheduled event to invoke the Lambda function.
- B. Review the AWS Trusted Advisor bucket permissions check and implement the recommended actions.
- C. Run a script that puts a private ACL on all of the objects in the bucket.
- D. Use the Block Public Access feature in Amazon S3 to set the IgnorePublicAcls option to TRUE on the bucket.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/access-control-block-public-access.html>

Question: 272

A large company recently experienced an unexpected increase in Amazon RDS and Amazon DynamoDB costs.

The company needs to increase visibility into details of AWS Billing and Cost Management. There are various accounts associated with AWS Organizations, including many development and production accounts. There is no consistent tagging strategy across the organization, but there are guidelines in place that require all infrastructure to be deployed using AWS CloudFormation with consistent tagging. Management requires cost center numbers and project ID numbers for all existing and future DynamoDB tables and RDS instances. Which strategy should the solutions architect provide to meet these requirements?

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

Answer: C

Explanation:

Using Tag Editor to remediate untagged resources is a Best Practice (Page 14 of AWS Tagging Best Practices WhitePaper). However, that is where answer A stops. It doesn't address the requirement of "Management requires cost center numbers and project ID number for all existing and future DynamoDB tables and RDS instances". That is where Answer C comes in and addresses that requirement with SCPs in the company's AWS Organization. AWS Tagging Best Practices - <https://d1.awsstatic.com/whitepapers/aws-tagging-best-practices.pdf>

Question: 273 A company is serving files to its customers through an SFTP server that is accessible over the internet. The SFTP server is running on a single Amazon EC2 instance with an Elastic IP address attached. Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication. The EC2 instance also has an attached security group that allows access from all customer IP addresses.

A solutions architect must implement a solution to improve availability, minimize the complexity of infrastructure management, and minimize the disruption to customers who access files. The solution must not change the way customers connect. Which solution will meet these requirements?

- A. Disassociate the Elastic IP address from the EC2 instance. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family server. Configure the Transfer Family server with a publicly accessible endpoint. Associate the SFTP Elastic IP address with the new endpoint. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.
- B. Disassociate the Elastic IP address from the EC2 instance. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family Server. Configure the Transfer Family server with a VPC-hosted, internet-facing endpoint. Associate the SFTP Elastic IP address with the new endpoint. Attach the security group with customer IP addresses to the new endpoint. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.
- C. Disassociate the Elastic IP address from the EC2 instance. Create a new Amazon Elastic File System (Amazon

EFS) file system to be used for SFTP file hosting. Create an AWS Fargate task definition to run an SFTP server Specify the EFS file system as a mount in the task definition Create a Fargate service by using the task definition, and place a Network Load Balancer (NLB) in front of the service. When configuring the service, attach the security group with customer IP addresses to the tasks that run the SFTP server Associate the Elastic IP address with the NLB Sync all files from the SFTP server to the S3 bucket.

D. Disassociate the Elastic IP address from the EC2 instance. Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hosting. Create a Network Load Balancer (NLB) with the Elastic IP address attached. Create an Auto Scaling group with EC2 instances that run an SFTP server. Define in the Auto Scaling group that instances that are launched should attach the new multi-attach EBS volume Configure the Auto Scaling group to automatically add instances behind the NLB. configure the Auto Scaling group to use the security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launches Sync all files from the SFTP server to the new multi-attach EBS volume.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/>

Question: 274

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 AWS Lambda function to copy objects from one S3 bucket to the other S3 bucket.

Answer: A,B,E

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/MultiRegionAccessPointRequestRouting.html>
<https://stackoverflow.com/questions/60947157/aws-s3-replication-without-versioning#:~:text=The%20automated%20Same%20Region%20Replication,is%20replicated%20between%20S3%20buckets.>

Question: 275

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

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

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

Answer: C,E

Explanation:

<https://aws.amazon.com/blogs/architecture/create-dynamic-contact-forms-for-s3-static-websites-using-aws-lambda-amazon-api-gateway-and-amazon-ses/>

Question: 276

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

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

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

Answer: A,C

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/privatelink-interface-endpoints.html#types-of-vpc-endpoints-for-s3>

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-bucket-access-direct-connect/>

Use a private IP address over Direct Connect (with an interface VPC endpoint)

To access Amazon S3 using a private IP address over Direct Connect, perform the following steps:

3. Create a private virtual interface for your connection.

5. Create an interface VPC endpoint for Amazon S3 in a VPC that is associated with the virtual private gateway.

The VGW must connect to a Direct Connect private virtual interface. This interface VPC endpoint resolves to a private IP address even if you enable a VPC endpoint for S3.

Question: 277

A company operates quick-service restaurants. The restaurants follow a predictable model with high sales traffic for 4 hours daily. Sales traffic is lower outside of those peak hours.

The point of sale and management platform is deployed in the AWS Cloud and has a backend that is based on Amazon DynamoDB. The database table uses provisioned throughput mode with 100,000 RCUs and 80,000 WCUs to match known peak resource consumption.

The company wants to reduce its DynamoDB cost and minimize the operational overhead for the IT staff. Which solution meets these requirements MOST cost-effectively?

- A. Reduce the provisioned RCUs and WCUs
- B. Change the DynamoDB table to use on-demand capacity.
- C. Enable DynamoDB auto scaling for the table
- D. Purchase 1-year reserved capacity that is sufficient to cover the peak load for 4 hours each day.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at-any-scale/> "As you can see, there are compelling reasons to use DynamoDB auto scaling with actively changing traffic. Auto scaling responds quickly and simplifies capacity management, which lowers costs by scaling your table's provisioned capacity and reducing operational overhead."

Question: 278

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: <https://aws.amazon.com/s3/features/access-points/>
<https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/>

Question: 279

A company is collecting a large amount of data from a fleet of IoT devices. Data is stored as Optimized Row Columnar (ORC) files in the Hadoop Distributed File System (HDFS) on a persistent Amazon EMR cluster. The company's data analytics team queries the data by using SQL in Apache Presto deployed on the same EMR cluster. Queries scan large amounts of data, always run for less than 15 minutes, and run only between 5 PM and 10 PM.

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

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

Answer: B

Explanation:

(<https://stackoverflow.com/questions/50250114/athena-vs-redshift-spectrum>)

Question: 280

An environmental company is deploying sensors in major cities throughout a country to measure air quality. The

sensors connect to AWS IoT Core to ingest timesheets data readings. The company stores the data in

Amazon DynamoDB

For business continuity the company must have the ability to ingest and store data in two AWS Regions Which solution will meet these requirements?

- A. Create an Amazon Route 53 alias failover routing policy with values for AWS IoT Core data endpoints in both Regions Migrate data to Amazon Aurora global tables
- B. Create a domain configuration for AWS IoT Core in each Region Create an Amazon Route 53 latency-based routing policy Use AWS IoT Core data endpoints in both Regions as values Migrate the data to Amazon MemoryDB for Redis and configure Cross-Region replication
- C. Create a domain configuration for AWS IoT Core in each. Region Create an AmazonRoute 53 health check that evaluates domain configuration health Create a failover routing policy with values for the domain name from the AWS IoT Core domain configurations Update the DynamoDB table to a global table
- D. Create an Amazon Route 53 latency-based routing policy. Use AWS IoT Core data endpoints in both Regions as values. Configure DynamoDB streams and Cross-Region data replication

Answer: C

Explanation:

<https://aws.amazon.com/solutions/implementations/disaster-recovery-for-aws-iot/>

Question: 281

A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration The company needs to store large, important documents within the application with the following requirements

- 1 The data must be highly durable and available
 2. The data must always be encrypted at rest and in transit.
 - 3 The encryption key must be managed by the company and rotated periodically
- Which of the following solutions should the solutions architect recommend?

- A. Deploy the storage gateway to AWS in file gateway mode Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes
- B. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.
- C. Use Amazon DynamoDB with SSL to connect to DynamoDB Use an AWS KMS key to encrypt DynamoDB objects at rest.
- D. Deploy instances with Amazon EBS volumes attached to store this data Use EBS volume encryption using an AWS KMS key to encrypt the data.

Answer: B

Explanation:

Question: 282

A solutions architect must update an application environment within AWS Elastic Beanstalk using a blue/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:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMEsSwap.html>

Question: 283

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:

because Amazon CloudFront considers the case of parameter names and values when caching based on query string parameters, thus inconsistent query strings may cause CloudFront to forward mixed-cased/misordered requests to the origin. Triggering a Lambda@Edge function based on a viewer request event to sort parameters by name and force them to be lowercase is the best choice.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/QueryStringParameters.html#query-string-parameters-optimizing-caching>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-cloudfront->

trigger-events.html

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-examples-normalize-query-string-parameters>

Question: 284

A solutions architect has implemented a SAML 2.0 federated identity solution with their company's on-premises identity provider (IdP) to authenticate users' access to the AWS environment. When the solutions architect tests authentication through the federated identity web portal, access to the AWS environment is granted.

However, when test users attempt to authenticate through the federated identity web portal, they are not able to access the AWS environment.

Which items should the solutions architect check to ensure identity federation is properly configured?

(Select THREE)

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

Answer: B,D,F

Explanation:

Question: 285

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 that have a statement of "Error Equals": ["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 that have a statement of "Error Equals": ["states.Runtime"] and "Next": "Email".

Answer: A,B,C

Explanation:

Create an Amazon Simple Notification Service (Amazon SNS) topic with a subscription of type "Email" that targets the team's mailing list. This will create a topic for sending notifications and add a subscription for the team's email list to that topic. C. Add a Catch field to all Task, Map, and Parallel states that have a statement of "ErrorEquals": ["States.ALL"] and "Next": "Email". This will ensure that any errors that occur in any of the steps in the workflow will trigger the "Email" task, which will forward the input arguments to the SNS topic created in step A. B. Create a task named "Email" that forwards the input arguments to the SNS topic. This will allow the company to send email notifications to the team's mailing list in case of any errors occurred in any step in the workflow.

Question: 286

A company is building an image service on the web that will allow users to upload and search random photos. At peak usage, up to 10,000 users worldwide will upload their images. The service will then overlay text on the uploaded images, which will then be published on the company website.

Which design should a solutions architect implement?

A. Store the uploaded images in Amazon Elastic File System (Amazon EFS). Send application log information about each image to Amazon CloudWatch Logs. Create a fleet of Amazon EC2 instances that use CloudWatch Logs to determine which images need to be processed. Place processed images in another directory in Amazon EFS. Enable Amazon CloudFront and configure the origin to be the one of the EC2 instances in the fleet.

B. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to Amazon Simple Notification Service (Amazon SNS). Create a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) to pull messages from Amazon SNS to process the images and place them in Amazon Elastic File System (Amazon EFS). Use Amazon CloudWatch metrics for the SNS message volume to scale out EC2 instances. Enable Amazon CloudFront and configure the origin to be the ALB in front of the EC2 instances.

C. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to the Amazon Simple Queue Service (Amazon SQS) queue. Create a fleet of Amazon EC2 instances to pull messages from the SQS queue to process the images and place them in another S3 bucket. Use Amazon CloudWatch metrics for queue depth to scale out EC2 instances. Enable Amazon CloudFront and configure the origin to be the S3 bucket that contains the processed images.

D. Store the uploaded images on a shared Amazon Elastic Block Store (Amazon EBS) volume. Create a fleet of Amazon EC2 Spot instances. Create an Amazon DynamoDB table that contains information about each uploaded image and whether it has been processed. Use an Amazon EventBridge rule to scale out EC2 instances. Enable Amazon CloudFront and configure the origin to reference an Elastic Load Balancer in front of the fleet of EC2 instances.

Answer: C

Explanation:

(Store the uploaded images in an S3 bucket and use S3 event notification with SQS queue) is the most suitable design. Amazon S3 provides highly scalable and durable storage for the uploaded images. Configuring S3 event notifications to send messages to an SQS queue allows for decoupling the processing of images from the upload process. A fleet of EC2 instances can pull messages from the SQS queue to process the images and store them in another S3 bucket. Scaling out the EC2 instances based on SQS queue depth using CloudWatch metrics ensures efficient utilization of resources. Enabling Amazon CloudFront with the origin set to the S3 bucket containing the processed images improves the global availability and performance of image delivery.

Question: 287

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 the source 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 point to the target DB instance endpoint.

Answer: A,C,E

Explanation:

Question: 288

A company is using AWS Organizations to manage multiple accounts. Due to regulatory requirements, the company wants to restrict specific member accounts to certain AWS Regions, where they are permitted to deploy resources. The resources in the accounts must be tagged, enforced based on a group standard and centrally managed with minimal configuration.

What should a solutions architect do to meet these requirements'?

- A. Create an AWS Config rule in the specific member accounts to limit Regions and apply a tag policy.
- B. From the AWS Billing and Cost Management console in the management account, disable Regions for the specific member accounts and apply a tag policy on the root.
- C. Associate the specific member accounts with the root. Apply a tag policy and an SCP using conditions to limit Regions.
- D. Associate the specific member accounts with a new OU. Apply a tag policy and an SCP using conditions to limit Regions.

Answer: D

Explanation:

<https://aws.amazon.com/es/blogs/mt/implement-aws-resource-tagging-strategy-using-aws-tag-policies-and-service-control-policies-scps/>

Question: 289

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 the data into an Amazon S3 bucket by using the S3 interface. 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 using the S3 Adapter for Snowball. Use AWS Application Migration Service to migrate the data from Amazon S3 to Aurora MySQL.

Answer: C

Explanation:

Question: 290

A company uses AWS CloudFormation to deploy applications within multiple VPCs that are all attached to a transit gateway. Each VPC that sends traffic to the public internet must send the traffic through a shared services VPC. Each subnet within a VPC uses the default VPC route table and the traffic is routed to the transit gateway. The transit gateway uses its default route table for any VPC attachment.

A security audit reveals that an Amazon EC2 instance that is deployed within a VPC can communicate with an EC2 instance that is deployed in any of the company's other VPCs. A solutions architect needs to limit the traffic between the VPCs. Each VPC must be able to communicate only with a predefined, limited set of authorized VPCs.

What should the solutions architect do to meet these requirements'?

- A. Update the network ACL of each subnet within a VPC to allow outbound traffic only to the authorized VPCs. Remove all deny rules except the default deny rule.
- B. Update all the security groups that are used within a VPC to deny outbound traffic to security groups that are used within the unauthorized VPCs.
- C. Create a dedicated transit gateway route table for each VPC attachment. Route traffic only to the authorized VPCs.
- D. Update the main route table of each VPC to route traffic only to the authorized VPCs through the transit gateway.

Answer: C

Explanation:

You can segment your network by creating multiple route tables in an AWS Transit Gateway and associate Amazon VPCs and VPNs to them. This will allow you to create isolated networks inside an AWS Transit Gateway similar to virtual routing and forwarding (VRFs) in traditional networks. The AWS Transit Gateway will have a default route table. The use of multiple route tables is optional.

Question: 291

A company has implemented an ordering system using an event-driven architecture. During initial testing, the system stopped processing orders. Further log analysis revealed that one order message in an Amazon Simple Queue Service (Amazon SQS) standard queue was causing an error on the backend and blocking all subsequent order messages. The visibility timeout of the queue is set to 30 seconds, and the backend processing timeout is set to 10 seconds. A solutions architect needs to analyze faulty order messages and ensure that the system continues to process subsequent messages.

Which step should the solutions architect take to meet these requirements?

- A. Increase the backend processing timeout to 30 seconds to match the visibility timeout.

- B. Reduce the visibility timeout of the queue to automatically remove the faulty message.
- C. Configure a new SQS FIFO queue as a dead-letter queue to isolate the faulty messages.
- D. Configure a new SQS standard queue as a dead-letter queue to isolate the faulty messages.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-dead-letter-queues.html>

Question: 292

A company has an IoT platform that runs in an on-premises environment. The platform consists of a server that connects to IoT devices by using the MQTT protocol. The platform collects telemetry data from the devices at least once every 5 minutes. The platform also stores device metadata in a MongoDB cluster.

An application that is installed on an on-premises machine runs periodic jobs to aggregate and transform the telemetry and device metadata. The application creates reports that users view by using another web

application that runs on the same on-premises machine. The periodic jobs take 120-600 seconds to run. However, the web application is always running.

The company is moving the platform to AWS and must reduce the operational overhead of the stack. Which combination of steps will meet these requirements with the LEAST operational overhead? (Select THREE.)

- A. Use AWS Lambda functions to connect to the IoT devices.
- B. Configure the IoT devices to publish to AWS IoT Core.
- C. Write the metadata to a self-managed MongoDB database on an Amazon EC2 instance.
- D. Write the metadata to Amazon DocumentDB (with MongoDB compatibility).
- E. Use AWS Step Functions state machines with AWS Lambda tasks to prepare the reports and to write the reports to Amazon S3. Use Amazon CloudFront with an S3 origin to serve the reports.
- F. Use an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with Amazon EC2 instances to prepare the reports. Use an ingress controller in the EKS cluster to serve the reports.

Answer: B,D,E

Explanation:

<https://aws.amazon.com/step-functions/use-cases/>

Question: 293

A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling Group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The

maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:

VPC CIDR: 10.0.0.0/23

AZ1 subnet CIDR: 10.0.0.0/24

AZ2 subnet CIDR: 10.0.1.0/24

Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime. Which solution will meet these requirements?

- A. Update the Auto Scaling group to use the AZ2 subnet only. Delete and re-create the AZ1 subnet using half the previous address space. Adjust the Auto Scaling group to also use the new AZ1 subnet. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only. Remove the current AZ2 subnet. Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.
- B. Terminate the EC2 instances in the AZ1 subnet. Delete and re-create the AZ1 subnet using half the address space. Update the Auto Scaling group to use this new subnet. Repeat this for the second AZ. Define a new subnet in AZ3; then update the Auto Scaling group to target all three new subnets.
- C. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ. Update the existing Auto Scaling group to target the new subnets in the new VPC.
- D. Update the Auto Scaling group to use the AZ2 subnet only. Update the AZ1 subnet to have half the previous address space. Adjust the Auto Scaling group to also use the AZ1 subnet again. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

Answer: A

Explanation:

<https://repost.aws/knowledge-center/vpc-ip-address-range>

Question: 294

A company hosts a web application on AWS in the us-east-1 Region. The application servers are distributed across three Availability Zones behind an Application Load Balancer. The database is hosted in a MySQL database on an Amazon EC2 instance. A solutions architect needs to design a Cross-Region data recovery solution using AWS services with an RTO of less than 5 minutes and an RPO of less than 1 minute. The solutions architect is deploying application servers in us-west-2, and has configured Amazon Route 53 health checks and DNS failover to us-west-2.

Which additional step should the solutions architect take?

- A. Migrate the database to an Amazon RDS for MySQL instance with a cross-Region read replica in us-west-2.
- B. Migrate the database to an Amazon Aurora global database with the primary in us-east-1 and the secondary in us-west-2.
- C. Migrate the database to an Amazon RDS for MySQL instance with a Multi-AZ deployment.

D. Create a MySQL standby database on an Amazon EC2 instance in us-west-2

Answer: B

Explanation:

<https://aws.amazon.com/rds/aurora/global-database/>

Question: 295

A company implements a containerized application by using Amazon Elastic Container Service (Amazon ECS) and Amazon API Gateway. The application data is stored in Amazon Aurora databases and Amazon DynamoDB databases. The company automates infrastructure provisioning by using AWS CloudFormation. The company automates application deployment by using AWS CodePipeline.

A solutions architect needs to implement a disaster recovery (DR) strategy that meets an RPO of 2 hours and an RTO of 4 hours.

Which solution will meet these requirements MOST cost-effectively'?

- A. Set up an Aurora global database and DynamoDB global tables to replicate the databases to a secondary AWS Region. In the primary Region and in the secondary Region, configure an API Gateway API with a Regional Endpoint. Implement Amazon CloudFront with origin failover to route traffic to the secondary Region during a DR scenario.
- B. Use AWS Database Migration Service (AWS DMS), Amazon EventBridge, and AWS Lambda to replicate the Aurora databases to a secondary AWS Region. Use DynamoDB Streams, EventBridge, and Lambda to replicate the DynamoDB databases to the secondary Region. In the primary Region and in the secondary Region, configure an API Gateway API with a Regional Endpoint. Implement Amazon Route 53 failover routing to switch traffic from the primary Region to the secondary Region.
- C. Use AWS Backup to create backups of the Aurora databases and the DynamoDB databases in a secondary AWS Region. In the primary Region and in the secondary Region, configure an API Gateway API with a Regional endpoint. Implement Amazon Route 53 failover routing to switch traffic from the primary Region to the secondary Region.
- D. Set up an Aurora global database and DynamoDB global tables to replicate the databases to a secondary AWS Region. In the primary Region and in the secondary Region, configure an API Gateway API with a Regional endpoint. Implement Amazon Route 53 failover routing to switch traffic from the primary Region to the secondary Region.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/database/cost-effective-disaster-recovery-for-amazon-aurora-databases-using-aws-backup/>

Question: 296

A company plans to deploy a new private intranet service on Amazon EC2 instances inside a VPC. An AWS Site-to-Site VPN connects the VPC to the company's on-premises network. The new service must communicate with existing on-premises services. The on-premises services are accessible through the use of hostnames that reside in the company example DNS zone. This DNS zone is wholly hosted on premises and is available only on the company's private network.

A solutions architect must ensure that the new service can resolve hostnames on the company example domain to integrate with existing services.

Which solution meets these requirements?

- A. Create an empty private zone in Amazon Route 53 for company example. Add an additional NS record to the company's on-premises company example zone that points to the authoritative name servers for the new private zone in Route 53.
- B. Turn on DNS hostnames for the VPC. Configure a new outbound endpoint with Amazon Route 53 Resolver. Create a Resolver rule to forward requests for company example to the on-premises name servers.
- C. Turn on DNS hostnames for the VPC. Configure a new inbound resolver endpoint with Amazon Route 53 Resolver. Configure the on-premises DNS server to forward requests for company example to the new resolver.
- D. Use AWS Systems Manager to configure a run document that will install a hosts file that contains any required hostnames. Use an Amazon EventBridge rule to run the document when an instance is entering the running state.

Answer: B

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver.html>

Question: 297

A company is using Amazon API Gateway to deploy a private REST API that will provide access to sensitive data. The API must be accessible only from an application that is deployed in a VPC. The company deploys the API successfully. However, the API is not accessible from an Amazon EC2 instance that is deployed in the VPC.

Which solution will provide connectivity between the EC2 instance and the API?

- A. Create an interface VPC endpoint for API Gateway. Attach an endpoint policy that allows `apigateway:*` actions. Disable private DNS naming for the VPC endpoint. Configure an API resource policy that allows access from the VPC. Use the VPC endpoint's DNS name to access the API.
- B. Create an interface VPC endpoint for API Gateway. Attach an endpoint policy that allows the `execute-api:Invoke` action. Enable private DNS naming for the VPC endpoint. Configure an API resource policy that allows access from the VPC endpoint. Use the API endpoint's DNS names to access the API. **Most Voted**

- C. Create a Network Load Balancer (NLB) and a VPC link. Configure private integration between API Gateway and the NLB. Use the API endpoint's DNS names to access the API.
- D. Create an Application Load Balancer (ALB) and a VPC Link. Configure private integration between API Gateway and the ALB. Use the ALB endpoint's DNS name to access the API.

Answer: B

Explanation:

According to the AWS documentation¹, to access a private API from a VPC, you need to do the following:

Create an interface VPC endpoint for API Gateway in your VPC. This creates a private connection between your VPC and API Gateway.

Attach an endpoint policy to the VPC endpoint that allows the execute-api:Invoke action for your private API. This grants permission to invoke your API from the VPC.

Enable private DNS naming for the VPC endpoint. This allows you to use the same DNS names for your private APIs as you would for public APIs.

Configure a resource policy for your private API that allows access from the VPC endpoint. This controls who can access your API and under what conditions.

Use the API endpoint's DNS names to access the API from your VPC. For example, `https://api-id.execute-api.region.amazonaws.com/stage`.

Question: 298

A company has developed a mobile game. The backend for the game runs on several virtual machines located in an on-premises data center. The business logic is exposed using a REST API with multiple functions. Player session data is stored in central file storage. Backend services use different API keys for throttling and to distinguish between live and test traffic.

The load on the game backend varies throughout the day. During peak hours, the server capacity is not sufficient. There are also latency issues when fetching player session data. Management has asked a solutions architect to present a cloud architecture that can handle the game's varying load and provide low-latency data access. The API model should not be changed. Which solution meets these requirements?

- A. Implement the REST API using a Network Load Balancer (NLB). Run the business logic on an Amazon EC2 instance behind the NLB. Store player session data in Amazon Aurora Serverless.
- B. Implement the REST API using an Application Load Balancer (ALB). Run the business logic in AWS Lambda. Store player session data in Amazon DynamoDB with on-demand capacity.
- C. Implement the REST API using Amazon API Gateway. Run the business logic in AWS Lambda. Store player session data in Amazon DynamoDB with on-demand capacity.
- D. Implement the REST API using AWS AppSync. Run the business logic in AWS Lambda. Store player session data in Amazon Aurora Serverless.

Answer: C

Explanation:

Question: 299

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

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

Answer: D

Explanation:

The correct answer is D. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

Option D meets the requirements of the scenario because it allows you to reduce operational overhead and support the product release by using the following AWS services and features: Amazon Elastic Kubernetes Service (Amazon EKS) is a fully managed service that allows you to run Kubernetes applications on AWS without needing to install, operate, or maintain your own Kubernetes control plane. You can use Amazon EKS to deploy your containerized application services on a Kubernetes cluster that is compatible with your existing tools and processes.

AWS Fargate is a serverless compute engine that eliminates the need to provision and manage servers for your containers. You can use AWS Fargate as the launch type for your Amazon EKS pods, which are the smallest deployable units of computing in Kubernetes. You can also enable auto scaling for your pods, which allows you to automatically adjust the number of pods based on the demand or custom metrics.

An Application Load Balancer (ALB) is a load balancer that distributes traffic across multiple targets in multiple Availability Zones using HTTP or HTTPS protocols. You can use an ALB to balance the load across your Amazon

EKS pods and provide high availability and fault tolerance for your application. Amazon RDS for PostgreSQL is a fully managed relational database service that supports the

PostgreSQL open source database engine. You can create additional read replicas for your DB instance, which are copies of your primary DB instance that can handle read-only queries and improve performance. You can also use read replicas to scale out beyond the capacity of a single DB instance for read-heavy workloads. Amazon Managed Streaming for Apache Kafka (Amazon MSK) is a fully managed service that makes it easy to build and run applications that use Apache Kafka to process streaming data. Apache Kafka is an open source platform for building real-time data pipelines and streaming applications. You can use Amazon MSK to create and manage a Kafka cluster that is highly available, secure, and compatible with your existing Kafka applications. You can also configure your application services to use the Amazon MSK cluster as a source or destination of streaming data.

Amazon S3 is an object storage service that offers high durability, availability, and scalability. You can store static content such as images, videos, or documents in Amazon S3 buckets, which are containers for objects. You can also serve static content directly from Amazon S3 using public URLs or presigned URLs.

Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds. You can use Amazon CloudFront to create a distribution that caches static content from your Amazon S3 bucket at edge locations closer to your users. This can improve the performance and user experience of your application.

Option A is incorrect because creating an EC2 Auto Scaling group behind an ALB would not reduce operational overhead as much as using AWS Fargate with Amazon EKS, as you would still need to manage EC2 instances for your containers. Creating additional read replicas for the DB instance would not provide high availability or fault tolerance in case of a failure of the primary DB instance, unlike deploying the DB instance in Multi-AZ mode.

Creating Amazon Kinesis data streams would not be compatible with your existing Apache Kafka applications, unlike using Amazon MSK.

Option B is incorrect because creating an EC2 Auto Scaling group behind an ALB would not reduce operational overhead as much as using AWS Fargate with Amazon EKS, as you would still need to manage EC2 instances for your containers. Creating Amazon Kinesis data streams would not be compatible with your existing Apache Kafka applications, unlike using Amazon MSK. Storing and serving static content directly from Amazon S3 would not provide optimal performance and user experience, unlike using Amazon CloudFront.

Option C is incorrect because deploying the application on a Kubernetes cluster created on the EC2 instances behind an ALB would not reduce operational overhead as much as using AWS Fargate with Amazon EKS, as you would still need to manage EC2 instances and Kubernetes control plane for your containers. Using Amazon API Gateway to interact with the application would add an unnecessary layer of complexity and cost to your architecture, as you would need to create and maintain an API gateway that proxies requests to your ALB.

Question: 300

A company wants to use AWS IAM Identity Center (AWS Single Sign-On) to manage employee access to AWS services. The company uses AWS Organizations to manage its AWS accounts.

Each employee has their own IAM user. Each IAM user is a member of at least one IAM group. Each IAM group has an attached policy that allows members to assume specific roles across the accounts. The roles contain appropriate policies for the expected activities of each group of users in each account. All relevant accounts exist inside a single OU.

The company has already created new users and groups in IAM Identity Center to match the permissions that exist in IAM.

How should the company use IAM Identity Center to implement the existing permissions?

- A. For each group, create policies in each account. Give the policies the same name in each account. Create a new permission set. Add the name of the new policies to the permission set. Assign user access to the AWS accounts in IAM Identity Center.
- B. For each group, create a new permission set. Attach the relevant existing IAM roles in each account to the permission set. Create a new customer managed policy that allows the group to assume the roles. Assign user access to the AWS accounts in IAM Identity Center.
- C. For each group, create a new permission set. Create policies in each account. Give each policy a unique name. Set the path of each policy to match the name of the permission set. Assign user access to the AWS accounts in IAM Identity Center.
- D. Add the OU to the accounts configuration in IAM Identity Center. For each group, create policies in each account. Create a new permission set. Add the new policies to the permission set as customer managed policies. Attach each new policy to the correct account in the account configuration in IAM Identity Center.

Answer: B

Explanation:

The correct answer is B. This option uses IAM Identity Center to create permission sets that map to the existing IAM roles in each account. This way, the company can leverage the existing policies and roles that are already configured for the expected activities of each group of users in each account. The company also needs to create a customer managed policy that allows the group to assume the roles and attach it to the permission set. This policy grants the necessary permissions for IAM Identity Center to assume the roles on behalf of the users. Finally, the company can assign user access to the AWS accounts in IAM Identity Center, which will automatically create IAM users and groups in each account based on the permission sets.

Option A is incorrect because it requires creating new policies in each account and giving them the same name. This is not necessary and adds complexity and overhead. The company can use the existing IAM roles and policies that are already configured for each account.

Option C is incorrect because it requires creating new policies in each account and giving them unique names. This is also not necessary and adds complexity and overhead. The company can use the existing IAM roles and policies that are already configured for each account.

Option D is incorrect because it requires adding the OU to the accounts configuration in IAM Identity Center. This is not supported by IAM Identity Center, which only allows adding individual accounts or all accounts in an organization.

Reference: AWS Single Sign-On Permission Sets

Question: 301 A financial services company sells its software-as-a-service (SaaS) platform for application compliance to large global banks. The SaaS platform runs on AWS and uses multiple AWS accounts that are managed in an organization in AWS Organizations. The SaaS platform uses many AWS resources globally.

For regulatory compliance, all API calls to AWS resources must be audited, tracked for changes, and stored in a durable and secure data store.

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

- A. Create a new AWS CloudTrail trail. Use an existing Amazon S3 bucket in the organization's management account to store the logs. Deploy the trail to all AWS Regions. Enable MFA delete and encryption on the S3 bucket.
- B. Create a new AWS CloudTrail trail in each member account of the organization. Create new Amazon S3 buckets to store the logs. Deploy the trail to all AWS Regions. Enable MFA delete and encryption on the S3 buckets.
- C. Create a new AWS CloudTrail trail in the organization's management account. Create a new Amazon S3 bucket with versioning turned on to store the logs. Deploy the trail for all accounts in the organization. Enable MFA delete and encryption on the S3 bucket.
- D. Create a new AWS CloudTrail trail in the organization's management account. Create a new Amazon S3 bucket to store the logs. Configure Amazon Simple Notification Service (Amazon SNS) to send log-file delivery notifications to an external management system that will track the logs. Enable MFA delete and encryption on the S3 bucket.

Answer: C

Explanation:

The correct answer is C. This option uses AWS CloudTrail to create a trail in the organization's management account that applies to all accounts in the organization. This way, the company can centrally manage and audit all API calls to AWS resources across multiple accounts and regions. The company also needs to create a new Amazon S3 bucket with versioning turned on to store the logs. Versioning helps protect against accidental or malicious deletion of log files by keeping multiple versions of each object in the bucket. The company also needs to enable MFA delete and encryption on the S3 bucket to further enhance the security and durability of the data store.

Option A is incorrect because it uses an existing S3 bucket in the organization's management account to store the logs. This may not be optimal for regulatory compliance, as the existing bucket may have different permissions, encryption settings, or lifecycle policies than a dedicated bucket for CloudTrail logs.

Option B is incorrect because it requires creating a new CloudTrail trail in each member account of the organization. This adds operational overhead and complexity, as the company would need to manage multiple trails and S3 buckets across multiple accounts and regions.

Option D is incorrect because it requires configuring Amazon SNS to send log-file delivery notifications to an external management system that will track the logs. This adds unnecessary complexity and cost, as CloudTrail already provides log-file integrity validation and log-file digest delivery features that can help verify the authenticity and integrity of log files.

Reference: Creating a Trail for an Organization

Question: 302 A company is migrating an application to the AWS Cloud. The application runs in an on-premises data center and writes thousands of images into a mounted NFS file system each night. After the company migrates the application, the company will host the application on an Amazon EC2 instance with a mounted Amazon

Elastic File System (Amazon EFS) file system.

The company has established an AWS Direct Connect connection to AWS. Before the migration cutover, a solutions architect must build a process that will replicate the newly created on-premises images to the EFS file system.

What is the MOST operationally efficient way to replicate the images?

- A. Configure a periodic process to run the aws s3 sync command from the on-premises file system to Amazon S3. Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system.
- B. Deploy an AWS Storage Gateway file gateway with an NFS mount point. Mount the file gateway file system on the on-premises server. Configure a process to periodically copy the images to the mount point.
- C. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. Send data over the Direct Connect connection to an S3 bucket by using public VIF. Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system.
- D. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. Send data over the Direct Connect connection to an AWS PrivateLink interface endpoint.

Answer: D

Explanation:

This option uses AWS DataSync to replicate the on-premises images to the EFS file system over the Direct Connect connection. AWS DataSync is a service that automates and accelerates data transfer between on-premises storage systems and AWS storage services. It can transfer data to and from Amazon EFS, Amazon FSx for Windows File Server, and Amazon S3. To use AWS DataSync, the company needs to deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. The agent connects to the AWS DataSync service endpoint in the AWS Region where the EFS file system is located. The company can use an AWS PrivateLink interface endpoint to connect to the service endpoint securely and privately over the Direct Connect connection. The company can then create a task in AWS DataSync that specifies the source location (the NFS file system), the destination location (the EFS file system), and the options for the data transfer (such as schedule, bandwidth limit, and verification). AWS DataSync will then perform the data transfer efficiently and securely, using encryption in transit and at rest.

Question: 303

A company runs its application on Amazon EC2 instances and AWS Lambda functions. The EC2 instances experience a continuous and stable load. The Lambda functions experience a varied and unpredictable load. The application includes a caching layer that uses an Amazon MemoryDB for Redis cluster.

A solutions architect must recommend a solution to minimize the company's overall monthly costs.

Which solution will meet these requirements?

- A. Purchase an EC2 Instance Savings Plan to cover the EC2 instances. Purchase a Compute Savings Plan for Lambda to cover the minimum expected consumption of the Lambda functions. Purchase reserved nodes to cover the MemoryDB cache nodes.
- B. Purchase a Compute Savings Plan to cover the EC2 instances. Purchase Lambda reserved concurrency to cover the expected Lambda usage. Purchase reserved nodes to cover the MemoryDB cache nodes.
- C. Purchase a Compute Savings Plan to cover the entire expected cost of the EC2 instances, Lambda

functions, and MemoryDB cache nodes.

D. Purchase a Compute Savings Plan to cover the EC2 instances and the MemoryDB cache nodes. Purchase Lambda reserved concurrency to cover the expected Lambda usage.

Answer: A

Explanation:

This option uses different types of savings plans and reserved nodes to minimize the company's overall monthly costs for running its application on EC2 instances, Lambda functions, and MemoryDB cache nodes. Savings plans are flexible pricing models that offer significant savings on AWS usage (up to 72%) in exchange for a commitment of a consistent amount of usage (measured in \$/hour) for a one-year or three-year term. There are two types of savings plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any compute usage across EC2 instances, Fargate containers, Lambda functions, SageMaker notebooks, and ECS tasks. EC2 Instance Savings Plans apply to a specific instance family within a region and provide more savings than Compute Savings Plans (up to 66% versus up to 54%). Reserved nodes are similar to savings plans but apply only to MemoryDB cache nodes. They offer up to 55% savings compared to on-demand pricing.

Question: 304

A company needs to monitor a growing number of Amazon S3 buckets across two AWS Regions. The company also needs to track the percentage of objects that are encrypted in Amazon S3. The company needs a dashboard to display this information for internal compliance teams.

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

A. Create a new S3 Storage Lens dashboard in each Region to track bucket and encryption metrics. Aggregate data from both Region dashboards into a single dashboard in Amazon QuickSight for the compliance teams.

B. Deploy an AWS Lambda function in each Region to list the number of buckets and the encryption status of objects. Store this data in Amazon S3. Use Amazon Athena queries to display the data on a custom dashboard in Amazon QuickSight for the compliance teams.

C. Use the S3 Storage Lens default dashboard to track bucket and encryption metrics. Give the compliance teams access to the dashboard directly in the S3 console.

D. Create an Amazon EventBridge rule to detect AWS Cloud Trail events for S3 object creation. Configure the rule to invoke an AWS Lambda function to record encryption metrics in Amazon DynamoDB. Use Amazon QuickSight to display the metrics in a dashboard for the compliance teams.

Answer: C

Explanation:

This option uses the S3 Storage Lens default dashboard to track bucket and encryption metrics across two AWS

Regions. S3 Storage Lens is a feature that provides organization-wide visibility into object storage usage and activity trends, and delivers actionable recommendations to improve cost efficiency and apply data protection best practices. S3 Storage Lens delivers more than 30 storage metrics, including metrics on encryption, replication, and data protection. The default dashboard provides a summary of the entire S3 usage and activity across all Regions and accounts in an organization. The company can give the compliance teams access to the dashboard directly in the S3 console, which requires the least operational overhead.

Question: 305

A company is planning to migrate an application to AWS. The application runs as a Docker container and uses an NFS version 4 file share.

A solutions architect must design a secure and scalable containerized solution that does not require provisioning or management of the underlying infrastructure.

Which solution will meet these requirements?

- A. Deploy the application containers by using Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type. Use Amazon Elastic File System (Amazon EFS) for shared storage. Reference the EFS file system ID, container mount point, and EFS authorization IAM role in the ECS task definition.
- B. Deploy the application containers by using Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type. Use Amazon FSx for Lustre for shared storage. Reference the FSx for Lustre file system ID, container mount point, and FSx for Lustre authorization IAM role in the ECS task definition.
- C. Deploy the application containers by using Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type and auto scaling turned on. Use Amazon Elastic File System (Amazon EFS) for shared storage. Mount the EFS file system on the ECS container instances. Add the EFS authorization IAM role to the EC2 instance profile.
- D. Deploy the application containers by using Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type and auto scaling turned on. Use Amazon Elastic Block Store (Amazon EBS) volumes with Multi-Attach enabled for shared storage. Attach the EBS volumes to ECS container instances. Add the EBS authorization IAM role to an EC2 instance profile.

Answer: A

Explanation:

This option uses Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type to deploy the application containers. Amazon ECS is a fully managed container orchestration service that allows running Docker containers on AWS at scale. Fargate is a serverless compute engine for containers that eliminates the need to provision or manage servers or clusters. With Fargate, the company only pays for the resources required to run its containers, which reduces costs and operational overhead. This option also uses Amazon Elastic File System (Amazon EFS) for shared storage. Amazon EFS is a fully managed file system that provides scalable, elastic, concurrent, and secure file storage for use with AWS cloud services. Amazon EFS supports NFS version 4 protocol, which is compatible with the application's requirements. To use Amazon EFS with Fargate

containers, the company needs to reference the EFS file system ID, container mount point, and EFS authorization IAM role in the ECS task definition.

Question: 306

A scientific company needs to process text and image data from an Amazon S3 bucket. The data is collected from several radar stations during a live, time-critical phase of a deep space mission. The radar stations upload the data to the source S3 bucket. The data is prefixed by radar station identification number.

The company created a destination S3 bucket in a second account. Data must be copied from the source S3 bucket to the destination S3 bucket to meet a compliance objective. The replication occurs through the use of an S3 replication rule to cover all objects in the source S3 bucket.

One specific radar station is identified as having the most accurate data.

a. Data replication at this radar station must be monitored for completion within 30 minutes after the radar station uploads the objects to the source S3 bucket.

What should a solutions architect do to meet these requirements?

A. Set up an AWS DataSync agent to replicate the prefixed data from the source S3 bucket to the destination S3 bucket. Select to use all available bandwidth on the task, and monitor the task to ensure that it is in the TRANSFERRING status. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.

B. In the second account, create another S3 bucket to receive data from the radar station with the most accurate data. Set up a new replication rule for this new S3 bucket to separate the replication from the other radar stations. Monitor the maximum replication time to the destination. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.

C. Enable Amazon S3 Transfer Acceleration on the source S3 bucket, and configure the radar station with the most accurate data to use the new endpoint. Monitor the S3 destination bucket's TotalRequestLatency metric. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.

D. Create a new S3 replication rule on the source S3 bucket that filters for the keys that use the prefix of the radar station with the most accurate data. Enable S3 Replication Time Control (S3 RTC). Monitor the maximum replication time to the destination. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-time-control.html>

Question: 307

A company is migrating a legacy application from an on-premises data center to AWS. The application consists of a single application server and a Microsoft SQL

Server database server. Each server is deployed on a VMware VM that consumes 500 TB of data across

multiple attached volumes.

The company has established a 10 Gbps AWS Direct Connect connection from the closest AWS Region to its on-premises data center. The Direct Connect connection is not currently in use by other services.

Which combination of steps should a solutions architect take to migrate the application with the LEAST amount of downtime? (Choose two.)

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

Answer: A,D

Explanation:

Question: 308

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 THREE.)

- A. Implement cross-account backup with AWS Backup vaults in designated non-production accounts.
- B. Add an SCP that restricts the modification of AWS Backup vaults.
- C. Implement AWS Backup Vault Lock in compliance mode.
- 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 nonproduction 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:

Question: 309

A company is expanding. The company plans to separate its resources into hundreds of different AWS accounts in multiple AWS Regions. A solutions architect must recommend a solution that denies access to any operations outside of specifically designated Regions.

Which solution will meet these requirements?

- A. Create IAM roles for each account. Create IAM policies with conditional allow permissions that include only approved Regions for the accounts.
- B. Create an organization in AWS Organizations. Create IAM users for each account. Attach a policy to each user to block access to Regions where an account cannot deploy infrastructure.
- C. Launch an AWS Control Tower landing zone. Create OUs and attach SCPs that deny access to run services outside of the approved Regions.
- D. Enable AWS Security Hub in each account. Create controls to specify the Regions where an account can deploy infrastructure.

Answer: C

Explanation:

Question: 310

A company is migrating its legacy .NET workload to AWS. The company has a containerized setup that includes a base container image. The base image is tens of gigabytes in size because of legacy libraries and other dependencies. The company has images for custom developed components that are dependent on the base image.

The company will use Amazon Elastic Container Registry (Amazon ECR) as part of its solution on AWS.

Which solution will provide the LOWEST container startup time on AWS?

- A. Use Amazon ECR to store the base image and the images for the custom developed components. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the workload.
- B. Use Amazon ECR to store the base image and the images for the custom developed components. Use AWS App Runner to run the workload.
- C. Use Amazon ECR to store the images for the custom developed components. Create an AMI that contains the base image. Use Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 instances that are

based on the AMI to run the workload

D. Use Amazon ECR to store the images for the custom developed components. Create an AMI that contains the base image. Use Amazon Elastic Kubernetes Service (Amazon EKS) on AWS Fargate with the AMI to run the workload.

Answer: C

Explanation:

Question: 311

A company has an application that uses an Amazon Aurora PostgreSQL DB cluster for the application's database. The DB cluster contains one small primary instance and three larger replica instances. The application runs on an AWS Lambda function. The application makes many short-lived connections to the database's replica instances to perform read-only operations.

During periods of high traffic, the application becomes unreliable and the database reports that too many connections are being established. The frequency of high-traffic periods is unpredictable.

Which solution will improve the reliability of the application?

- A. Use Amazon RDS Proxy to create a proxy for the DB cluster. Configure a read-only endpoint for the proxy. Update the Lambda function to connect to the proxy endpoint.
- B. Increase the max_connections setting on the DB cluster's parameter group. Reboot all the instances in the DB cluster. Update the Lambda function to connect to the DB cluster endpoint.
- C. Configure instance scaling for the DB cluster to occur when the DatabaseConnections metric is close to the max_connections setting. Update the Lambda function to connect to the Aurora reader endpoint.
- D. Use Amazon RDS Proxy to create a proxy for the DB cluster. Configure a read-only endpoint for the Aurora Data API on the proxy. Update the Lambda function to connect to the proxy endpoint.

Answer: A

Explanation:

Question: 312

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

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

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

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

Answer: B

Explanation:

Question: 313

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

EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.

What should a solutions architect do to meet these requirements?

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

Answer: D

Explanation:

Question: 314

A company is running a workload that consists of thousands of Amazon EC2 instances. The workload is running in a VPC that contains several public subnets and private subnets. The public subnets have a route for 0.0.0.0/0 to an existing internet gateway. The private subnets have a route for 0.0.0.0/0 to an existing NAT gateway. A solutions architect needs to migrate the entire fleet of EC2 instances to use IPv6. The EC2 instances that are in private subnets must not be accessible from the public internet. What should the solutions architect do to meet these requirements?

- A. Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnets. Update all the VPC route tables, and add a route for ::/0 to the internet gateway.
- B. Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Update the VPC route tables for all private subnets, and add a route for ::/0 to the NAT gateway.
- C. Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Create an egress-only internet gateway. Update the VPC route tables for all private subnets, and add a route for ::/0 to the egress-only internet gateway.
- D. Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnets. Create a new NAT gateway, and enable IPv6 support. Update the VPC route tables for all private subnets, and add a route for ::/0 to the IPv6-enabled NAT gateway.

Answer: C

Explanation:

Question: 315

A Solutions Architect wants to make sure that only AWS users or roles with suitable permissions can access a new Amazon API Gateway endpoint. The Solutions Architect wants an end-to-end view of each request to analyze the latency of the request and create service maps.

How can the Solutions Architect design the API Gateway access control and perform request inspections?

- A. For the API Gateway method, set the authorization to AWS_IAM. Then, give the IAM user or role execute-api:Invoke permission on the REST API resource. Enable the API caller to sign requests with AWS Signature when accessing the endpoint. Use AWS X-Ray to trace and analyze user requests to API Gateway.
- B. For the API Gateway resource, set CORS to enabled and only return the company's domain in Access-Control-Allow-Origin headers. Then, give the IAM user or role execute-api:Invoke permission on the REST API resource. Use Amazon CloudWatch to trace and analyze user requests to API Gateway.
- C. Create an AWS Lambda function as the custom authorizer, ask the API client to pass the key and secret

when making the call, and then use Lambda to validate the key/secret pair against the IAM system. Use AWS X-Ray to trace and analyze user requests to API Gateway.

D. Create a client certificate for API Gateway. Distribute the certificate to the AWS users and roles that need to access the endpoint. Enable the API caller to pass the client certificate when accessing the endpoint. Use Amazon CloudWatch to trace and analyze user requests to API Gateway.

Answer: A

Explanation:

Question: 316

A North American company with headquarters on the East Coast is deploying a new web application running on Amazon EC2 in the us-east-1 Region. The application should dynamically scale to meet user demand and maintain resiliency. Additionally, the application must have disaster recover capabilities in an active-passive configuration with the us-west-1 Region.

Which steps should a solutions architect take after creating a VPC in the us-east-1 Region?

A. Create a VPC in the us-west-1 Region. Use inter-Region VPC peering to connect both VPCs. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region. Deploy EC2 instances across multiple AZs in each Region as part of an Auto Scaling group spanning both VPCs and served by the ALB.

B. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the ALB. Deploy the same solution to the us-west-1 Region. Create an Amazon Route 53 record set with a failover routing policy and health checks enabled to provide high availability across both Regions.

C. Create a VPC in the us-west-1 Region. Use inter-Region VPC peering to connect both VPCs. Deploy an Application Load Balancer (ALB) that spans both VPCs. Deploy EC2 instances across multiple Availability Zones as part of an Auto Scaling group in each VPC served by the ALB. Create an Amazon Route 53 record that points to the ALB.

D. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the ALB. Deploy the same solution to the us-west-1 Region. Create separate Amazon Route 53 records in each Region that point to the ALB in the Region. Use Route 53 health checks to provide high availability across both Regions.

Answer: B

Explanation:

Question: 317

A company hosts a data-processing application on Amazon EC2 instances. The application polls an Amazon

Elastic File System (Amazon EFS) file system for newly uploaded files. When a new file is detected, the application extracts data from the file and runs logic to select a Docker container image to process the file. The application starts the appropriate container image and passes the file location as a parameter.

The data processing that the container performs can take up to 2 hours. When the processing is complete, the code that runs inside the container writes the file back to Amazon EFS and exits.

The company needs to refactor the application to eliminate the EC2 instances that are running the containers

Which solution will meet these requirements?

- A. Create an Amazon Elastic Container Service (Amazon ECS) cluster. Configure the processing to run as AWS Fargate tasks. Extract the container selection logic to run as an Amazon EventBridge rule that starts the appropriate Fargate task. Configure the EventBridge rule to run when files are added to the EFS file system.
- B. Create an Amazon Elastic Container Service (Amazon ECS) cluster. Configure the processing to run as AWS Fargate tasks. Update and containerize the container selection logic to run as a Fargate service that starts the appropriate Fargate task. Configure an EFS event notification to invoke the Fargate service when files are added to the EFS file system.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster. Configure the processing to run as AWS Fargate tasks. Extract the container selection logic to run as an AWS Lambda function that starts the appropriate Fargate task. Migrate the storage of file uploads to an Amazon S3 bucket. Update the processing code to use Amazon S3. Configure an S3 event notification to invoke the Lambda function when objects are created.
- D. Create AWS Lambda container images for the processing. Configure Lambda functions to use the container images. Extract the container selection logic to run as a decision Lambda function that invokes the appropriate Lambda processing function. Migrate the storage of file uploads to an Amazon S3 bucket. Update the processing code to use Amazon S3. Configure an S3 event notification to invoke the decision Lambda function when objects are created.

Answer: D

Explanation:

Question: 318

A company runs a web application on AWS. The web application delivers static content from an Amazon S3 bucket that is behind an Amazon CloudFront distribution. The application serves dynamic content by using an Application Load Balancer (ALB) that distributes requests to a fleet of Amazon EC2 instances in Auto Scaling groups. The application uses a domain name setup in Amazon Route 53.

Some users reported occasional issues when the users attempted to access the website during peak hours. An operations team found that the ALB sometimes returned HTTP 503 Service Unavailable errors. The company wants to display a custom error message page when these errors occur. The page should be displayed immediately for this error code.

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

- A. Set up a Route 53 failover routing policy. Configure a health check to determine the status of the ALB endpoint and to fail over to the failover S3 bucket endpoint.
- B. Create a second CloudFront distribution and an S3 static website to host the custom error page. Set up a Route 53 failover routing policy. Use an active-passive configuration between the two distributions.
- C. Create a CloudFront origin group that has two origins. Set the ALB endpoint as the primary origin. For the secondary origin, set an S3 bucket that is configured to host a static website. Set up origin failover for the CloudFront distribution. Update the S3 static website to incorporate the custom error page.
- D. Create a CloudFront function that validates each HTTP response code that the ALB returns. Create an S3 static website in an S3 bucket. Upload the custom error page to the S3 bucket as a failover. Update the function to read the S3 bucket and to serve the error page to the end users.

Answer: B

Explanation:

Question: 319

A software as a service (SaaS) company uses AWS to host a service that is powered by AWS PrivateLink. The service consists of proprietary software that runs on three Amazon EC2 instances behind a Network Load Balancer (NLB). The instances are in private subnets in multiple Availability Zones in the eu-west-2 Region. All the company's customers are in eu-west-2.

However, the company now acquires a new customer in the us-east-1 Region. The company creates a new VPC and new subnets in us-east-1. The company establishes inter-Region VPC peering between the VPCs in the two Regions.

The company wants to give the new customer access to the SaaS service, but the company does not want to immediately deploy new EC2 resources in us-east-1

Which solution will meet these requirements?

- A. Configure a PrivateLink endpoint service in us-east-1 to use the existing NLB that is in eu-west-2. Grant specific AWS accounts access to connect to the SaaS service.
- B. Create an NLB in us-east-1. Create an IP target group that uses the IP addresses of the company's instances in eu-west-2 that host the SaaS service. Configure a PrivateLink endpoint service that uses the NLB that is in us-east-1. Grant specific AWS accounts access to connect to the SaaS service.
- C. Create an Application Load Balancer (ALB) in front of the EC2 instances in eu-west-2. Create an NLB in us-east-1. Associate the NLB that is in us-east-1 with an ALB target group that uses the ALB that is in eu-west-2. Configure a PrivateLink endpoint service that uses the NLB that is in us-east-1. Grant specific AWS accounts access to connect to the SaaS service.
- D. Use AWS Resource Access Manager (AWS RAM) to share the EC2 instances that are in eu-west-2. In us-east-1, create an NLB and an instance target group that includes the shared EC2 instances from eu-west-2. Configure

a PrivateLink endpoint service that uses the NL B that is in us-east-1. Grant specific AWS accounts access to connect to the SaaS service.

Answer: B

Explanation:

Question: 320

A company's CISO has asked a Solutions Architect to re-engineer the company's current CI/CD practices to make sure patch deployments to its applications can happen as quickly as possible with minimal downtime if vulnerabilities are discovered. The company must also be able to quickly roll back a change in case of errors.

The web application is deployed in a fleet of Amazon EC2 instances behind an Application Load Balancer. The company is currently using GitHub to host the application source code, and has configured an AWS CodeBuild project to build the application. The company also intends to use AWS CodePipeline to trigger builds from GitHub commits using the existing CodeBuild project.

What CI/CD configuration meets all of the requirements?

- A. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for in-place deployment. Monitor the newly deployed code, and, if there are any issues, push another code update.
- B. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for blue/green deployments. Monitor the newly deployed code, and, if there are any issues, trigger a manual rollback using CodeDeploy.
- C. Configure CodePipeline with a deploy stage using AWS CloudFormation to create a pipeline for test and production stacks. Monitor the newly deployed code, and, if there are any issues, push another code update.
- D. Configure the CodePipeline with a deploy stage using AWS OpsWorks and in-place deployments. Monitor the newly deployed code, and, if there are any issues, push another code update.

Answer: B

Explanation:

Question: 321

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:

Question: 322

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:

Question: 323

A company has multiple AWS accounts. The company recently had a security audit that revealed many unencrypted Amazon Elastic Block Store (Amazon EBS) volumes attached to Amazon EC2 instances. A solutions architect must encrypt the unencrypted volumes and ensure that unencrypted volumes will be detected automatically in the future. Additionally, the company wants a solution that can centrally manage multiple AWS accounts with a focus on compliance and security.

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

- A. Create an organization in AWS Organizations. Set up AWS Control Tower, and turn on the strongly recommended guardrails. Join all accounts to the organization. Categorize the AWS accounts into OUs.
- B. Use the AWS CLI to list all the unencrypted volumes in all the AWS accounts. Run a script to encrypt all the unencrypted volumes in place.
- C. Create a snapshot of each unencrypted volume. Create a new encrypted volume from the unencrypted snapshot. Detach the existing volume, and replace it with the encrypted volume.
- D. Create an organization in AWS Organizations. Set up AWS Control Tower, and turn on the mandatory

guardrails. Join all accounts to the organization. Categorize the AWS accounts into OUs.

E. Turn on AWS CloudTrail. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to detect and automatically encrypt unencrypted volumes.

Answer: A,C

Explanation:

(<https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-guardrails.html>)

Question: 324

A company wants to use Amazon S3 to back up its on-premises file storage solution. The company's on-premises file storage solution supports NFS, and the company wants its new solution to support NFS. The company wants to archive the backup files after 5 days. If the company needs archived files for disaster recovery, the company is willing to wait a few days for the retrieval of those files. Which solution meets these requirements MOST cost-effectively?

- A. Deploy an AWS Storage Gateway files gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the file gateway. Create an S3 Lifecycle rule to move the file to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.
- B. Deploy an AWS Storage Gateway volume gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the volume gateway. Create an S3 Lifecycle rule to move the files to S3 Glacier Deep Archive after 5 days.
- C. Deploy an AWS Storage Gateway tape gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the tape gateway. Create an S3 Lifecycle rule to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.
- D. Deploy an AWS Storage Gateway file gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the tape gateway. Create an S3 Lifecycle rule to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.
- E. Deploy an AWS Storage Gateway file gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the file gateway. Create an S3 Lifecycle rule to move the files to S3 Glacier Deep Archive after 5 days.

Answer: E

Explanation:

Question: 325

A research company is running daily simulations in the AWS Cloud to meet high demand. The simulations run on several hundred Amazon EC2 instances that are based on Amazon Linux 2. Occasionally, a simulation gets stuck and requires a cloud operations engineer to solve the problem by connecting to an EC2 instance through SSH.

Company policy states that no EC2 instance can use the same SSH key and that all connections must be logged in AWS CloudTrail.

How can a solutions architect meet these requirements?

A. Launch new EC2 instances, and generate an individual SSH key for each instance. Store the SSH key in AWS Secrets Manager. Create a new IAM policy, and attach it to the engineers' IAM role with an Allow statement for the GetSecretValue action. Instruct the engineers to fetch the SSH key from Secrets Manager when they connect through any SSH client.

B. Create an AWS Systems Manager document to run commands on EC2 instances to set a new unique SSH key. Create a new IAM policy, and attach it to the engineers' IAM role with an Allow statement to run Systems Manager documents. Instruct the engineers to run the document to set an SSH key and to connect through any SSH client.

C. Launch new EC2 instances without setting up any SSH key for the instances. Set up EC2 Instance Connect on each instance. Create a new IAM policy, and attach it to the engineers' IAM role with an Allow statement for the SendSSHPublicKey action. Instruct the engineers to connect to the instance by using a browser-based SSH client from the EC2 console.

D. Set up AWS Secrets Manager to store the EC2 SSH key. Create a new AWS Lambda function to create a new SSH key and to call AWS Systems Manager Session Manager to set the SSH key on the EC2 instance. Configure Secrets Manager to use the Lambda function for automatic rotation once daily. Instruct the engineers to fetch the SSH key from Secrets Manager when they connect through any SSH client.

Answer: C

Explanation:

Question: 326

A financial services company has an asset management product that thousands of customers use around the world. The customers provide feedback about the product through surveys. The company is building a new analytical solution that runs on Amazon EMR to analyze the data from these surveys. The following user personas need to access the analytical solution to perform different actions:

- Administrator: Provisions the EMR cluster for the analytics team based on the team's requirements
- Data engineer: Runs ETL scripts to process, transform, and enrich the datasets
- Data analyst: Runs SQL and Hive queries on the data

A solutions architect must ensure that all the user personas have least privilege access to only the resources that they need. The user personas must be able to launch only applications that are approved and authorized. The solution also must ensure tagging for all resources that the user personas create.

Which solution will meet these requirements?

A. Create IAM roles for each user persona. Attach identity-based policies to define which actions the user who

assumes the role can perform. Create an AWSConfig rule to check for noncompliant resources. Configure the rule to notify the administrator to remediate the noncompliant resources.

B. Set up Kerberos-based authentication for EMR clusters upon launch. Specify a Kerberos security configuration along with cluster-specific Kerberos options.

C. Use AWS Service Catalog to control the Amazon EMR versions available for deployment, the cluster configuration, and the permissions for each user persona.

D. Launch the EMR cluster by using AWS CloudFormation. Attach resource-based policies to the EMR cluster during cluster creation. Create an AWS Config rule to check for noncompliant clusters and noncompliant Amazon S3 buckets. Configure the rule to notify the administrator to remediate the noncompliant resources.

Answer: C

Explanation:

Question: 327

A company operates a fleet of servers on premises and operates a fleet of Amazon EC2 instances in its organization in AWS Organizations. The company's AWS accounts contain hundreds of VPCs. The company wants to connect its AWS accounts to its on-premises network. AWS Site-to-Site VPN connections are already established to a single AWS account. The company wants to control which VPCs can communicate with other VPCs.

Which combination of steps will achieve this level of control with the LEAST operational effort? (Choose three.)

A. Create a transit gateway in an AWS account. Share the transit gateway across accounts by using AWS Resource Access Manager (AWS RAM).

B. Configure attachments to all VPCs and VPNs.

C. Set up transit gateway route tables. Associate the VPCs and VPNs with the route tables.

D. Configure VPC peering between the VPCs.

E. Configure attachments between the VPCs and VPNs.

F. Set up route tables on the VPCs and VPNs.

Answer: A,B,C

Explanation:

Question: 328

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? (Select 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:

Question: 329

A company has hundreds of AWS accounts. The company uses an organization in AWS Organizations to manage all the accounts. The company has turned on all features.

A finance team has allocated a daily budget for AWS costs. The finance team must receive an email notification if the organization's AWS costs exceed 80% of the allocated budget. A solutions architect needs to implement a solution to track the costs and deliver the notifications.

Which solution will meet these requirements?

- A. In the organization's management account, use AWS Budgets to create a budget that has a daily period. Add an alert threshold and set the value to 80%. Use Amazon Simple Notification Service (Amazon SNS) to notify the finance team.
- B. In the organization's management account, set up the organizational view feature for AWS Trusted Advisor. Create an organizational view report for cost optimization. Set an alert threshold of 80%. Configure notification preferences. Add the email addresses of the finance team.
- C. Register the organization with AWS Control Tower. Activate the optional cost control (guardrail). Set a control (guardrail) parameter of 80%. Configure control (guardrail) notification preferences. Use Amazon Simple Notification Service (Amazon SNS) to notify the finance team.
- D. Configure the member accounts to save a daily AWS Cost and Usage Report to an Amazon S3 bucket in the organization's management account. Use Amazon EventBridge to schedule a daily Amazon Athena query to calculate the organization's costs. Configure Athena to send an Amazon CloudWatch alert if the total costs are more than 80% of the allocated budget. Use Amazon Simple Notification Service (Amazon SNS) to notify the finance team.

Answer: A

Explanation:

Question: 330

A company is migrating mobile banking applications to run on Amazon EC2 instances in a VPC. Backend service applications run in an on-premises data center. The data center has an AWS Direct Connect connection into AWS. The applications that run in the VPC need to resolve DNS requests to an on-premises Active Directory domain that runs in the data center.

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

- A. Provision a set of EC2 instances across two Availability Zones in the VPC as caching DNS servers to resolve DNS queries from the application servers within the VPC.
- B. Provision an Amazon Route 53 private hosted zone. Configure NS records that point to on-premises DNS servers.
- C. Create DNS endpoints by using Amazon Route 53 Resolver. Add conditional forwarding rules to resolve DNS namespaces between the on-premises data center and the VPC.
- D. Provision a new Active Directory domain controller in the VPC with a bidirectional trust between this new domain and the on-premises Active Directory domain.

Answer: C

Explanation:

Question: 331

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:

Question: 332

A company needs to aggregate Amazon CloudWatch logs from its AWS accounts into one central logging account. The collected logs must remain in the AWS Region of creation. The central logging account will then process the logs, normalize the logs into standard output format, and stream the output logs to a security tool for more processing.

A solutions architect must design a solution that can handle a large volume of logging data that needs to be ingested. Less logging will occur outside normal business hours than during normal business hours. The logging solution must scale with the anticipated load. The solutions architect has decided to use an AWS Control Tower design to handle the multi-account logging process.

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

- A. Create a destination Amazon Kinesis data stream in the central logging account.
- B. Create a destination Amazon Simple Queue Service (Amazon SQS) queue in the central logging account.
- C. Create an IAM role that grants Amazon CloudWatch Logs the permission to add data to the Amazon Kinesis data stream. Create a trust policy. Specify the trust policy in the IAM role. In each member account, create a subscription filter for each log group to send data to the Kinesis data stream.
- D. Create an IAM role that grants Amazon CloudWatch Logs the permission to add data to the Amazon Simple Queue Service (Amazon SQS) queue. Create a trust policy. Specify the trust policy in the IAM role. In each member account, create a single subscription filter for all log groups to send data to the SQS queue.
- E. Create an AWS Lambda function. Program the Lambda function to normalize the logs in the central logging account and to write the logs to the security tool.
- F. Create an AWS Lambda function. Program the Lambda function to normalize the logs in the member accounts and to write the logs to the security tool.

Answer: A,C,E

Explanation:

Question: 333

A company has an organization in AWS Organizations that includes a separate AWS account for each of the company's departments. Application teams from different departments develop and deploy solutions independently.

The company wants to reduce compute costs and manage costs appropriately across departments. The company also wants to improve visibility into billing for individual departments. The company does not want to lose operational flexibility when the company selects compute resources.

Which solution will meet these requirements?

- www.atmicnetworks.com
- A. Use AWS Budgets for each department. Use Tag Editor to apply tags to appropriate resources. Purchase EC2 Instance Savings Plans.
 - B. Configure AWS Organizations to use consolidated billing. Implement a tagging strategy that identifies departments. Use SCPs to apply tags to appropriate resources. Purchase EC2 Instance Savings Plans.
 - C. Configure AWS Organizations to use consolidated billing. Implement a tagging strategy that identifies departments. Use Tag Editor to apply tags to appropriate resources. Purchase Compute Savings Plans.
 - D. Use AWS Budgets for each department. Use SCPs to apply tags to appropriate resources. Purchase Compute Savings Plans.

Answer: C

Explanation:

Question: 334

A large payroll company recently merged with a small staffing company. The unified company now has multiple business units, each with its own existing AWS account.

A solutions architect must ensure that the company can centrally manage the billing and access policies for all the AWS accounts. The solutions architect configures AWS Organizations by sending an invitation to all member accounts of the company from a centralized management account.

What should the solutions architect do next to meet these requirements?

- A. Create the OrganizationAccountAccess IAM group in each member account. Include the necessary IAM roles for each administrator.
- B. Create the OrganizationAccountAccessPolicy IAM policy in each member account. Connect the member accounts to the management account by using cross-account access.
- C. Create the OrganizationAccountAccessRole IAM role in each member account. Grant permission to the management account to assume the IAM role.
- D. Create the OrganizationAccountAccessRole IAM role in the management account. Attach the AdministratorAccess AWS managed policy to the IAM role. Assign the IAM role to the administrators in each member account.

Answer: D

Explanation:

Question: 335

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 repository and 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 repository endpoint.
- 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 to use the CodeArtifact repository. Create a VPC endpoint for CodeArtifact.

Answer: D

Explanation:

Question: 336

A solutions architect works for a government agency that has strict disaster recovery requirements. All Amazon Elastic Block Store (Amazon EBS) snapshots are required to be saved in at least two additional AWS Regions. The agency also is required to maintain the lowest possible operational overhead.

Which solution meets these requirements?

- A. Configure a policy in Amazon Data Lifecycle Manager (Amazon DLM) to run once daily to copy the EBS snapshots to the additional Regions.
- B. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to copy the EBS snapshots to the additional Regions.
- C. Set up AWS Backup to create the EBS snapshots. Configure Amazon S3 cross-Region replication to copy

the EBS snapshots to the additional Regions.

D. Schedule Amazon EC2 Image Builder to run once daily to create an AMI and copy the AMI to the additional Regions

Answer: A

Explanation:

Question: 337

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 VPC. Ensure that the EC2 instance is joined to the domain. Use the bastion host to access the target instances through RDP.
- 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 Microsoft Active Directory AD Connector. Configure permission sets against user groups for access to AWS Systems Manager. Use Systems Manager Fleet Manager to access the target instances through RDP.
- C. Implement a VPN between the on-premises environment and the target VPC. Ensure that the target instances are joined to the on-premises Active Directory domain over the VPN connection. Configure RDP access through the VPN. Connect from the company's network to the target instances.
- 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: D

Explanation:

Question: 338

A company is running an application in the AWS Cloud. The application uses AWS Lambda functions and Amazon Elastic Container Service (Amazon ECS) containers that run with AWS Fargate technology as its primary compute. The load on the application is irregular. The application experiences long periods of no usage, followed by sudden and significant increases and decreases in traffic. The application is write-heavy and stores

data in an Amazon Aurora MySQL database. The database runs on an Amazon RDS memory optimized DB instance that is not able to handle the load.

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

- A. Add additional read replicas to the database. Purchase Instance Savings Plans and RDS Reserved Instances.
- B. Migrate the database to an Aurora multi-master DB cluster. Purchase Instance Savings Plans.
- C. Migrate the database to an Aurora global database. Purchase Compute Savings Plans and RDS Reserved Instances.
- D. Migrate the database to Aurora Serverless v1. Purchase Compute Savings Plans.

Answer: D

Explanation:

Question: 339

A company is using AWS CodePipeline for the CI/CD of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS

CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts.

As the application has become more complex, recent resource changes in the CloudFormation templates have caused unplanned downtime.

How should a solutions architect improve the CI/CD pipeline to reduce the likelihood that changes in the templates will cause downtime?

- A. Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployments. Write test plans for a testing team to execute in a non-production environment before approving the change for production.
- B. Implement automated testing using AWS CodeBuild in a test environment. Use CloudFormation change sets to evaluate changes before deployment. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.
- C. Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correct. Adapt the deployment code to check for error conditions and generate notifications on errors. Deploy to a test environment and execute a manual test plan before approving the change for production.
- D. Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment scripts. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.

Answer: B

Explanation:

Question: 340

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

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

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

Which combination of steps should a solutions architect recommend to avoid a recurrence of this issue?

(Choose two.)

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

Answer: C,D

Explanation:

Question: 341

A solutions architect is reviewing an application's resilience before launch. The application runs on an Amazon EC2 instance that is deployed in a private subnet of a VPC.

The EC2 instance is provisioned by an Auto Scaling group that has a minimum capacity of 1 and a maximum capacity of 1. The application stores data on an Amazon RDS for MySQL DB instance. The VPC has subnets configured in three Availability Zones and is configured with a single NAT gateway.

The solutions architect needs to recommend a solution to ensure that the application will operate across multiple Availability Zones.

Which solution will meet this requirement?

- A. Deploy an additional NAT gateway in the other Availability Zones. Update the route tables with appropriate routes. Modify the RDS for MySQL DB instance to a Multi-AZ configuration. Configure the Auto Scaling group to launch instances across Availability Zones. Set the minimum capacity and maximum capacity of the Auto Scaling group to 3.
- B. Replace the NAT gateway with a virtual private gateway. Replace the RDS for MySQL DB instance with an Amazon Aurora MySQL DB cluster. Configure the Auto Scaling group to launch instances across all subnets in the VPC. Set the minimum capacity and maximum capacity of the Auto Scaling group to 3.
- C. Replace the NAT gateway with a NAT instance. Migrate the RDS for MySQL DB instance to an RDS for PostgreSQL DB instance. Launch a new EC2 instance in the other Availability Zones.
- D. Deploy an additional NAT gateway in the other Availability Zones. Update the route tables with appropriate routes. Modify the RDS for MySQL DB instance to turn on automatic backups and retain the backups for 7 days. Configure the Auto Scaling group to launch instances across all subnets in the VPC. Keep the minimum capacity and the maximum capacity of the Auto Scaling group at 1.

Answer: A

Explanation:

Question: 342

A company hosts an application on AWS. The application reads and writes objects that are stored in a single Amazon S3 bucket. The company must modify the application to deploy the application in two AWS Regions.

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

- A. Set up an Amazon CloudFront distribution with the S3 bucket as an origin. Deploy the application to a second Region. Modify the application to use the CloudFront distribution. Use AWS Global Accelerator to access the data in the S3 bucket.
- B. Create a new S3 bucket in a second Region. Set up bidirectional S3 Cross-Region Replication (CRR) between the original S3 bucket and the new S3 bucket. Configure an S3 Multi-Region Access Point that uses both S3 buckets. Deploy a modified application to both Regions.
- C. Create a new S3 bucket in a second Region. Deploy the application in the second Region. Configure the application to use the new S3 bucket. Set up S3 Cross-Region Replication (CRR) from the original S3 bucket to the new S3 bucket.
- D. Set up an S3 gateway endpoint with the S3 bucket as an origin. Deploy the application to a second Region. Modify the application to use the new S3 gateway endpoint. Use S3 Intelligent-Tiering on the S3 bucket.

Answer: B

Explanation:

Question: 343

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 in 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 times for a global audience? (Choose 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 in 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 the web and application tiers are each in Auto Scaling groups.

Answer: D,E

Explanation:

Question: 344

A company needs to implement a disaster recovery (DR) plan for a web application. The application runs in a single AWS Region.

The application uses microservices that run in containers. The containers are hosted on AWS Fargate in Amazon Elastic Container Service (Amazon ECS). The application has an Amazon RDS for MySQL DB instance as its data layer and uses Amazon Route 53 for DNS resolution. An Amazon CloudWatch alarm invokes an Amazon EventBridge rule if the application experiences a failure.

A solutions architect must design a DR solution to provide application recovery to a separate Region. The solution must minimize the time that is necessary to recover from a failure.

Which solution will meet these requirements?

- A. Set up a second ECS cluster and ECS service on Fargate in the separate Region. Create an AWS Lambda function to perform the following actions: take a snapshot of the RDS DB instance. copy the snapshot to the separate Region. create a new RDS DB instance from the snapshot, and update Route 53 to route traffic to the second ECS cluster. Update the EventBridge rule to add a target that will invoke the Lambda function.
- B. Create an AWS Lambda function that creates a second ECS cluster and ECS service in the separate Region. Configure the Lambda function to perform the following actions: take a snapshot of the RDS DB instance, copy the snapshot to the separate Region. create a new RDS DB instance from the snapshot. and update Route 53 to route traffic to the second ECS cluster. Update the EventBridge rule to add a target that will invoke the Lambda function.
- C. Set up a second ECS cluster and ECS service on Fargate in the separate Region. Create a cross-Region read replica of the RDS DB instance in the separate Region. Create an AWS Lambda function to promote the read replica to the primary database. Configure the Lambda function to update Route 53 to route traffic to the second ECS cluster. Update the EventBridge rule to add a target that will invoke the Lambda function.
- D. Set up a second ECS cluster and ECS service on Fargate in the separate Region. Take a snapshot of the RDS DB instance. Convert the snapshot to an Amazon DynamoDB global table. Create an AWS Lambda function to update Route 53 to route traffic to the second ECS cluster. Update the EventBridge rule to add a target that will invoke the Lambda function.

Answer: C

Explanation:

This option uses a cross-Region read replica of the RDS DB instance to provide a standby database in the separate Region. A cross-Region read replica is a copy of the primary database that is updated asynchronously using the native replication features of the database engine. It provides enhanced availability, scalability, and performance for read-heavy workloads. It also enables fast recovery from a regional outage by promoting the read replica to a standalone database. To use a cross-Region read replica, the company needs to set up a second ECS cluster and ECS service on Fargate in the separate Region. The company also needs to create an AWS Lambda function to promote the read replica to the primary database and update Route 53 to route traffic to the second ECS cluster. The company can then update the EventBridge rule to add a target that will invoke the Lambda function in case of a failure.

Question: 345

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

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

- A. Create a new security group and attach it to the CloudFront distribution. Update the ALB security group ingress to allow access only from the CloudFront security group.
- B. Update ALB security group ingress to allow access only from the CloudFront managed prefix list.
- C. Create a VPC interface endpoint for Elastic Load Balancing. Update the ALB scheme from internet-facing to internal_
- D. Extract CloudFront IPS from the AWS provided ip-ranges.json document. Update ALB security group ingress to allow access only from CloudFront IPs.

Answer: B

Explanation:

The CloudFront managed prefix list contains the IP ranges for all CloudFront edge locations. By updating the ALB security group ingress to allow access only from this prefix list, the web application will be accessible only by using the CloudFront endpoint. This solution requires the least amount of effort compared to the other options, which involve creating new resources or updating existing ones. This solution also avoids hard-coding IP addresses, which can change over time.

Reference: section "Security and Compliance"

Question: 346

A live-events company is designing a scaling solution for its ticket application on AWS. The application has high peaks of utilization during sale events. Each sale event is a one-time event that is scheduled.

The application runs on Amazon EC2 instances that are in an Auto Scaling group. The application uses PostgreSQL for the database layer.

The company needs a scaling solution to maximize availability during the sale events.

Which solution will meet these requirements?

- A. Use a predictive scaling policy for the EC2 instances. Host the database on an Amazon Aurora PostgreSQL Serverless v2 Multi-AZ DB instance with automatically scaling read replicas. Create an AWS Step Functions state machine to run parallel AWS Lambda functions to pre-warm the database before a sale event. Create an Amazon EventBridge rule to invoke the state machine.
- B. Use a scheduled scaling policy for the EC2 instances. Host the database on an Amazon RDS for PostgreSQL Multi-AZ DB instance with automatically scaling read replicas. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger read replica before a sale event. Fail over to the larger read replica. Create another EventBridge rule that invokes another Lambda function to scale down the read replica after the sale event.
- C. Use a predictive scaling policy for the EC2 instances. Host the database on an Amazon RDS for PostgreSQL Multi-AZ DB instance with automatically scaling read replica. Create an AWS Step Functions state machine to run parallel AWS Lambda functions to pre-warm the database before a sale event. Create an Amazon

EventBridge rule to invoke the state machine.

D. Use a scheduled scaling policy for the EC2 instances. Host the database on an Amazon Aurora PostgreSQL Multi-AZ DB cluster. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger Aurora Replica before a sale event. Fail over to the larger Aurora Replica. Create another EventBridge rule that invokes another Lambda function to scale down the Aurora Replica after the sale event.

Answer: D

Explanation:

The correct answer is D. Use a scheduled scaling policy for the EC2 instances. Host the database on an Amazon Aurora PostgreSQL Multi-AZ DB cluster. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger Aurora Replica before a sale event. Fail over to the larger Aurora Replica. Create another EventBridge rule that invokes another Lambda function to scale down the Aurora Replica after the sale event.

This solution will meet the requirements of maximizing availability during the sale events. A scheduled scaling policy for the EC2 instances will allow the application to scale up and down according to the predefined schedule of the sale events. Hosting the database on an Amazon Aurora PostgreSQL Multi-AZ DB cluster will provide high availability and durability, as well as compatibility with PostgreSQL. Creating an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger Aurora Replica before a sale event will ensure that the database can handle the increased read traffic during the peak periods. Failing over to the larger Aurora Replica will make it the primary instance, which will also improve the write performance of the database. Creating another EventBridge rule that invokes another Lambda function to scale down the Aurora Replica after the sale event will reduce the cost and resources of the database.

Reference: [3], section "Scaling Amazon Aurora MySQL and PostgreSQL with Aurora Auto Scaling"

Question: 347

A company's solutions architect is evaluating an AWS workload that was deployed several years ago. The application tier is stateless and runs on a single large Amazon EC2 instance that was launched from an AMI. The application stores data in a MySQL database that runs on a single EC2 instance.

The CPU utilization on the application server EC2 instance often reaches 100% and causes the application to stop responding. The company manually installs patches on the instances. Patching has caused downtime in the past. The company needs to make the application highly available.

Which solution will meet these requirements with the LEAST development time?

A. Move the application tier to AWS Lambda functions in the existing VPC. Create an Application Load Balancer to distribute traffic across the Lambda functions. Use Amazon GuardDuty to scan the Lambda functions. Migrate the database to Amazon DocumentDB (with MongoDB compatibility).

B. Change the EC2 instance type to a smaller Graviton powered instance type. Use the existing AMI to create a launch template for an Auto Scaling group. Create an Application Load Balancer to distribute traffic across the instances in the Auto Scaling group. Set the Auto Scaling group to scale based on CPU utilization. Migrate

the database to Amazon DynamoDB.

C. Move the application tier to containers by using Docker. Run the containers on Amazon Elastic Container Service (Amazon ECS) with EC2 instances. Create an Application Load Balancer to distribute traffic across the ECS cluster. Configure the ECS cluster to scale based on CPU utilization. **Migrate the database to Amazon Neptune.**

D. Create a new AMI that is configured with AWS Systems Manager Agent (SSM Agent). Use the new AMI to create a launch template for an Auto Scaling group. Use smaller instances in the Auto Scaling group. Create an Application Load Balancer to distribute traffic across the instances in the Auto Scaling group. Set the Auto Scaling group to scale based on CPU utilization. **Migrate the database to Amazon Aurora MySQL.**

Answer: D

Explanation:

This solution will meet the requirements of making the application highly available with the least development time. Creating a new AMI that is configured with SSM Agent will enable the company to use AWS Systems Manager to manage and patch the EC2 instances automatically, reducing downtime and human errors. Using a launch template for an Auto Scaling group will allow the company to launch multiple instances of the same configuration and scale them up or down based on demand. Using smaller instances in the Auto Scaling group will reduce the cost and improve the performance of the application tier. Creating an Application Load Balancer to distribute traffic across the instances in the Auto Scaling group will increase the availability and fault tolerance of the application tier. Migrating the database to Amazon Aurora MySQL will provide a fully managed, compatible, and scalable relational database service that can handle high throughput and concurrent connections.

Question: 348

A solutions architect needs to migrate an on-premises legacy application to AWS. The application runs on two servers behind a load balancer. The application requires a license file that is associated with the MAC address of the server's network adapter. It takes the software vendor 12 hours to send new license files. The application also uses configuration files with a static IP address to access a database host names are not supported.

Given these requirements, which combination of steps should be taken to implement highly available architecture for the application servers in AWS? (Select TWO.)

- A. Create a pool of ENIs. Request license files from the vendor for the pool, and store the license files in Amazon S3. Create a bootstrap automation script to download a license file and attach the corresponding ENI to an Amazon EC2 instance.
- B. Create a pool of ENIs. Request license files from the vendor for the pool, store the license files on an Amazon EC2 instance. Create an AMI from the instance and use this AMI for all future EC2 instances.
- C. Create a bootstrap automation script to request a new license file from the vendor. When the response is received, apply the license file to an Amazon EC2 instance.
- D. Edit the bootstrap automation script to read the database server IP address from the AWS Systems Manager Parameter Store, and inject the value into the local configuration files.

E. Edit an Amazon EC2 instance to include the database server IP address in the configuration files and re-create the AMI to use for all future EC2 instances.

Answer: A,D

Explanation:

This solution will meet the requirements of implementing a highly available architecture for the application servers in AWS. Creating a pool of ENIs will allow the application servers to have consistent MAC addresses, which are needed for the license files. Requesting license files from the vendor for the pool and storing them in Amazon S3 will ensure that the license files are available and secure. Creating a bootstrap automation script to download a license file and attach the corresponding ENI to an EC2 instance will automate the process of launching new application servers with valid licenses. Editing the bootstrap automation script to read the database server IP address from the AWS Systems Manager Parameter Store and inject the value into the local configuration files will enable the application servers to access the database without hard-coding the IP address in the configuration files. This will also allow changing the database server IP address without modifying the configuration files on each application server.

Question: 349

A company migrated an application to the AWS Cloud. The application runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). Application data is stored in a MySQL database that runs on an additional EC2 instance. The application's use of the database is read-heavy.

The loads static content from Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instance. The static content is updated frequently and must be copied to each EBS volume.

The load on the application changes throughout the day. During peak hours, the application cannot handle all the incoming requests. Trace data shows that the database cannot handle the read load during peak hours.

Which solution will improve the reliability of the application?

A. Migrate the application to a set of AWS Lambda functions. Set the Lambda functions as targets for the ALB. Create a new single EBS volume for the static content. Configure the Lambda functions to read from the new EBS volume. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB cluster.

B. Migrate the application to a set of AWS Step Functions state machines. Set the state machines as targets for the ALB. Create an Amazon Elastic File System (Amazon EFS) file system for the static content. Configure the state machines to read from the EFS file system. Migrate the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance.

C. Containerize the application. Migrate the application to an Amazon Elastic Container Service (Amazon ECS) Cluster. Use the AWS Fargate launch type for the tasks that host the application. Create a new single EBS volume the static content. Mount the new EBS volume on the ECS duster. Configure AWS Application Auto Scaling on ECS cluster. Set the ECS service as a target for the ALB. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB cluster.

D. Containerize the application. Migrate the application to an Amazon Elastic Container Service (Amazon ECS) cluster. Use the AWS Fargate launch type for the tasks that host the application. Create an Amazon Elastic File System (Amazon EFS) file system for the static content. Mount the EFS file system to each container. Configure AWS Application Auto Scaling on the ECS cluster. Set the ECS service as a target for the ALB. Migrate the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance.

Answer: D

Explanation:

This solution will improve the reliability of the application by addressing the issues of scalability, availability, and performance. Containerizing the application will make it easier to deploy and manage on AWS. Migrating the application to an Amazon ECS cluster will allow the application to run on a fully managed container orchestration service. Using the AWS Fargate launch type for the tasks that host the application will enable the application to run on serverless compute engines that are automatically provisioned and scaled by AWS. Creating an Amazon EFS file system for the static content will provide a scalable and shared storage solution that can be accessed by multiple containers. Mounting the EFS file system to each container will eliminate the need to copy the static content to each EBS volume and ensure that the content is always up to date. Configuring AWS Application Auto Scaling on the ECS cluster will enable the application to scale up and down based on demand or a predefined schedule. Setting the ECS service as a target for the ALB will distribute the incoming requests across multiple tasks in the ECS cluster and improve the availability and fault tolerance of the application. Migrating the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance will provide a fully managed, compatible, and scalable relational database service that can handle high throughput and concurrent connections. Using a reader DB instance will offload some of the read load from the primary DB instance and improve the performance of the database.

Question: 350

A company has AWS accounts that are in an organization in AWS Organizations. The company wants to track Amazon EC2 usage as a metric. The company's architecture team must receive a daily alert if the EC2 usage is more than 10% higher than the average EC2 usage from the last 30 days.

Which solution will meet these requirements?

A. Configure AWS Budgets in the organization's management account. Specify a usage type of EC2 running hours. Specify a daily period. Set the budget amount to be 10% more than the reported average usage for the last 30 days from AWS Cost Explorer. Configure an alert to notify the architecture team if the usage threshold is met.

B. Configure AWS Cost Anomaly Detection in the organization's management account. Configure a monitor type of AWS Service. Apply a filter of Amazon EC2. Configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days.

C. Enable AWS Trusted Advisor in the organization's management account. Configure a cost optimization advisory alert to notify the architecture team if the EC2 usage is 10% more than the reported average usage

for the last 30 days.

D. Configure Amazon Detective in the organization's management account. Configure an EC2 usage anomaly alert to notify the architecture team if Detective identifies a usage anomaly of more than 10%.

Answer: B

Explanation:

AWS Cost Anomaly Detection is a feature of the AWS Cost Management suite that leverages machine learning to enable continuous monitoring of your AWS costs and usage, allowing you to identify unexpected and abnormal spending¹. You can create cost monitors that evaluate specific AWS services, member accounts, cost allocation tags, or cost categories based on your AWS account structure². You can also configure alert subscriptions that notify you when a cost monitor detects an anomaly that meets your threshold². In this case, you can create a cost monitor with a monitor type of AWS Service and apply a filter of Amazon EC2 to track the EC2 usage as a metric. You can then configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days, which is the anomaly detection period used by AWS Cost Anomaly Detection³.

Question: 351

A company is running an application on premises. The application uses a set of web servers that host a static React-based single-page application (SPA), a Node.js API, and a MySQL database server. The database is read intensive. The company will need to expand the database's storage at an unpredictable rate.

The company must migrate the application to AWS. The company also must modernize the architecture to reduce infrastructure management and increase scalability.

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

A. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon RDS for MySQL. Use AWS Application Migration Service to migrate the web application to a fleet of Amazon EC2 instances behind an Elastic Load Balancing (ELB) load balancer. Use a Spot Fleet with a request type of request to host the API.

B. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon Aurora MySQL. Copy the web files to an Amazon S3 bucket and set up web hosting. Copy the API code to AWS Lambda functions. Configure Amazon API Gateway to point to the Lambda functions.

C. Use AWS Database Migration Service (AWS DMS) to migrate the database to a MySQL database that runs on Amazon EC2 instances. Use AWS DataSync to migrate the web files and API files to an Amazon FSx for Windows File Server file system. Set up a fleet of EC2 instances in an Auto Scaling group as web servers. Mount the FSx for Windows File Server file system.

D. Use AWS Application Migration Service to migrate the database to Amazon EC2 instances. Copy the web files to containers that run on Amazon Elastic Kubernetes Service (Amazon EKS). Set up an Elastic Load Balancing (ELB) load balancer for the EC2 instances and EKS containers. Copy the API code to AWS Lambda functions. Configure Amazon API Gateway to point to the Lambda functions.

Answer: B

Explanation:

Question: 352

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 OUs 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:

This option will meet the requirements of preventing any new or existing EC2 instances in the OU's accounts from gaining a public IP address. An SCP is a policy that you can attach to an OU or an account in AWS Organizations to define the maximum permissions for the entities in that OU or account. By creating an SCP that denies the ec2:RunInstances and ec2:AssociateAddress actions when the value of the aws:RequestTag/aws:PublicIp condition key is true, you can prevent any user or role in the OU from launching instances that have a public IP address or attaching a public IP address to existing instances. This will effectively enforce a security best practice and reduce the risk of unauthorized access to your EC2 instances.

Question: 353

A company is planning to migrate its on-premises VMware cluster of 120 VMS to AWS. The VMS have many different operating systems and many custom software packages installed. The company also has an on-premises NFS server that is 10 TB in size. The company has set up a 10 Gbps AWS Direct Connect connection to AWS for the migration

Which solution will complete the migration to AWS in the LEAST amount of time?

- A. Export the on-premises VMS and copy them to an Amazon S3 bucket. Use VM Import/Export to create AMIS

- from the VM images that are stored in Amazon S3. Order an AWS Snowball Edge device. Copy the NFS server data to the device. Restore the NFS server data to an Amazon EC2 instance that has NFS configured.
- B. Configure AWS Application Migration Service with a connection to the VMware cluster. Create a replication job for the VMS. Create an Amazon Elastic File System (Amazon EFS) file system. Configure AWS DataSync to copy the NFS server data to the EFS file system over the Direct Connect connection.
- C. Recreate the VMS on AWS as Amazon EC2 instances. Install all the required software packages. Create an Amazon FSx for Lustre file system. Configure AWS DataSync to copy the NFS server data to the FSx for Lustre file system over the Direct Connect connection.
- D. Order two AWS Snowball Edge devices. Copy the VMS and the NFS server data to the devices. Run VM Import/Export after the data from the devices is loaded to an Amazon S3 bucket. Create an Amazon Elastic File System (Amazon EFS) file system. Copy the NFS server data from Amazon S3 to the EFS file system.

Answer: B

Explanation:

This option will complete the migration to AWS in the least amount of time because it uses two AWS services that are designed to simplify and accelerate data transfers and migrations. AWS Application Migration Service (AWS MGN) is a highly automated lift-and-shift solution that helps you migrate applications from any source infrastructure that runs supported operating systems to AWS¹. It replicates your source servers into your AWS account and automatically converts and launches them on AWS so you can quickly benefit from the cloud¹. You can use AWS MGN to migrate your on-premises VMware VMs to AWS by configuring a connection to your VMware cluster and creating a replication job for the VMs². This process will minimize the time-intensive, error-prone manual processes of exporting and importing VM images.

AWS DataSync is an online data movement and discovery service that simplifies and accelerates data migrations to AWS and helps you move data quickly and securely between on-premises storage, edge locations, other cloud providers, and AWS Storage³. It can transfer data between Network File System (NFS) shares, Server Message Block (SMB) shares, Hadoop Distributed File Systems (HDFS), self-managed object storage, AWS Snowcone, Amazon Simple Storage Service (Amazon S3) buckets, Amazon Elastic File System (Amazon EFS) file systems, Amazon FSx for Windows File Server file systems, Amazon FSx for Lustre file systems, Amazon FSx for OpenZFS file systems, and Amazon FSx for NetApp ONTAP file systems³. You can use AWS DataSync to copy your on-premises NFS server data to an Amazon EFS file system over the Direct Connect connection⁴. This process will leverage the high bandwidth and low latency of Direct Connect and the encryption and data integrity validation of DataSync.

Question: 354

An ecommerce company runs an application on AWS. The application has an Amazon API Gateway API that invokes an AWS Lambda function. The data is stored in an Amazon RDS for PostgreSQL DB instance.

During the company's most recent flash sale, a sudden increase in API calls negatively affected the application's performance. A solutions architect reviewed the Amazon CloudWatch metrics during that time and noticed a significant increase in Lambda invocations and database connections. The CPU utilization also was high on the DB instance.

What should the solutions architect recommend to optimize the application's performance?

- A. Increase the memory of the Lambda function. Modify the Lambda function to close the database connections when the data is retrieved.
- B. Add an Amazon ElastiCache for Redis cluster to store the frequently accessed data from the RDS database.
- C. Create an RDS proxy by using the Lambda console. Modify the Lambda function to use the proxy endpoint.
- D. Modify the Lambda function to connect to the database outside of the function's handler. Check for an existing database connection before creating a new connection.

Answer: C

Explanation:

This option will optimize the application's performance by reducing the overhead of opening and closing database connections for each Lambda invocation. An RDS proxy is a fully managed database proxy for Amazon RDS that makes applications more scalable, more resilient to database failures, and more secure¹. It allows applications to pool and share connections established with the database, improving database efficiency and application scalability¹. By creating an RDS proxy by using the Lambda console, you can easily configure your Lambda function to use the proxy endpoint instead of the direct database endpoint². This will enable your Lambda function to reuse existing connections from the proxy's connection pool, reducing the latency and CPU utilization caused by establishing new connections for each invocation. It will also prevent connection saturation or exhaustion on the database, which can degrade performance or cause errors³.

Question: 355

A company is building an application that will run on an AWS Lambda function. Hundreds of customers will use the application. The company wants to give each customer a quota of requests for a specific time period. The quotas must match customer usage patterns. Some customers must receive a higher quota for a shorter time period.

Which solution will meet these requirements?

- A. Create an Amazon API Gateway REST API with a proxy integration to invoke the Lambda function. For each customer, configure an API Gateway usage plan that includes an appropriate request quota. Create an API key from the usage plan for each user that the customer needs.
- B. Create an Amazon API Gateway HTTP API with a proxy integration to invoke the Lambda function. For each customer, configure an API Gateway usage plan that includes an appropriate request quota. Configure route-level throttling for each usage plan. Create an API key from the usage plan for each user that the customer needs.
- C. Create a Lambda function alias for each customer. Include a concurrency limit with an appropriate request quota. Create a Lambda function URL for each function alias. Share the Lambda function URL for each alias

with the relevant customer.

D. Create an Application Load Balancer (ALB) in a VPC. Configure the Lambda function as a target for the ALB. Configure an AWS WAF web ACL for the ALB. For each customer, configure a rate-based rule that includes an appropriate request quota.

Answer: A

Explanation:

The correct answer is A.

A . This solution meets the requirements because it allows the company to create different usage plans for each customer, with different request quotas and time periods. The usage plans can be associated with API keys, which can be distributed to the users of each customer. The API Gateway REST API can invoke the Lambda function using a proxy integration, which passes the request data to the function as input and returns the function output as the response. This solution is scalable, secure, and cost-effective¹²

B . This solution is incorrect because API Gateway HTTP APIs do not support usage plans or API keys. These features are only available for REST APIs³

C . This solution is incorrect because it does not provide a way to enforce request quotas for each customer. Lambda function aliases can be used to create different versions of the function, but they do not have any quota mechanism. Moreover, this solution exposes the Lambda function URLs directly to the customers, which is not secure or recommended⁴

D . This solution is incorrect because it does not provide a way to differentiate between customers or users. AWS WAF rate-based rules can be used to limit requests based on IP addresses, but they do not support any other criteria such as user agents or headers. Moreover, this solution adds unnecessary complexity and cost by using an ALB and a VPC⁵⁶

Reference:

1:Creating and using usage plans with API keys - Amazon API Gateway²:Set up a proxy integration with a Lambda proxy integration - Amazon API Gateway³:Choose between HTTP APIs and REST APIs - Amazon API Gateway⁴:Using AWS Lambda aliases - AWS Lambda⁵:Rate-based rule statement - AWS

WAF, AWS Firewall Manager, and AWS Shield Advanced⁶:Lambda functions as targets for Application Load Balancers - Elastic Load Balancing

Question: 356

A live-events company is designing a scaling solution for its ticket application on AWS. The application has high peaks of utilization during sale events. Each sale event is a one-time event that is scheduled. The application runs on Amazon EC2 instances that are in an Auto Scaling group.

The application uses PostgreSQL for the database layer.

The company needs a scaling solution to maximize availability during the sale events.

Which solution will meet these requirements?

- A. Use a predictive scaling policy for the EC2 instances. Host the database on an Amazon Aurora PostgreSQL Serverless v2 Multi-AZ DB instance with automatically scaling read replicas. Create an AWS Step Functions state machine to run parallel AWS Lambda functions to pre-warm the database before a sale event. Create an Amazon EventBridge rule to invoke the state machine.
- B. Use a scheduled scaling policy for the EC2 instances. Host the database on an Amazon RDS for PostgreSQL Multi-AZ DB instance with automatically scaling read replicas. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger read replica before a sale event. Fail over to the larger read replica. Create another EventBridge rule that invokes another Lambda function to scale down the read replica after the sale event.
- C. Use a predictive scaling policy for the EC2 instances. Host the database on an Amazon RDS for PostgreSQL Multi-AZ DB instance with automatically scaling read replicas. Create an AWS Step Functions state machine to run parallel AWS Lambda functions to pre-warm the database before a sale event. Create an Amazon EventBridge rule to invoke the state machine.
- D. Use a scheduled scaling policy for the EC2 instances. Host the database on an Amazon Aurora PostgreSQL Multi-AZ DB cluster. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger Aurora Replica before a sale event. Fail over to the larger Aurora Replica. Create another EventBridge rule that invokes another Lambda function to scale down the Aurora Replica after the sale event.

Answer: D

Explanation:

The correct answer is D.

D . This solution meets the requirements because it uses a scheduled scaling policy for the EC2 instances, which can adjust the capacity according to the known sale events. It also uses Amazon Aurora PostgreSQL Multi-AZ DB cluster, which provides high availability and durability for the database. It uses Amazon EventBridge rules and AWS Lambda functions to create a larger Aurora Replica before a sale event and fail over to it, which can improve the performance and handle the increased traffic. It also uses another EventBridge rule and Lambda function to scale down the Aurora

Replica after the sale event, which can save costs¹²³

A . This solution is incorrect because it uses predictive scaling policy for the EC2 instances, which is not suitable for one-time events that are scheduled. Predictive scaling is based on historical data and machine learning, which may not accurately forecast the demand for sale events. It also uses Amazon Aurora PostgreSQL Serverless v2 Multi-AZ DB instance, which does not support read replicas. The use of AWS Step Functions state machine and Lambda functions to pre-warm the database is unnecessary and adds complexity⁴⁵

B . This solution is incorrect because it uses Amazon RDS for PostgreSQL Multi-AZ DB instance with automatically scaling read replicas, which may not provide enough performance improvement for the sale events. The use of EventBridge rules and Lambda functions to create a larger read replica and fail over to it is risky and may cause downtime or data loss. The use of another EventBridge rule and Lambda function to scale down the read replica is also risky and may cause inconsistency or data loss⁶⁷

C . This solution is incorrect because it uses predictive scaling policy for the EC2 instances, which is not suitable for one-time events that are scheduled. Predictive scaling is based on historical data and machine learning, which may not accurately forecast the demand for sale events. The use of AWS Step Functions state machine and Lambda functions to pre-warm the database is unnecessary and adds complexity⁴⁵

Reference:

1:Scheduled scaling for Amazon EC2 Auto Scaling2:Amazon Aurora PostgreSQL features3:Amazon EventBridge rules4:Predictive scaling for Amazon EC2 Auto Scaling5:Amazon Aurora Serverless v26:Multi-AZ DB instance deployments - Amazon Relational Database Service7:Working with PostgreSQL read replicas - Amazon Relational Database Service

Question: 357

A company is deploying a new API to AWS. The API uses Amazon API Gateway with a Regional API endpoint and an AWS Lambda function for hosting. The API retrieves data from an external vendor API, stores data in an Amazon DynamoDB global table, and retrieves data from the DynamoDB global table. The API key for the vendor's API is stored in AWS Secrets Manager and is encrypted with a customer managed key in AWS Key Management Service (AWS KMS). The company has deployed its own API into a single AWS Region.

A solutions architect needs to change the API components of the company's API to ensure that the components can run across multiple Regions in an active-active configuration.

Which combination of changes will meet this requirement with the LEAST operational overhead? (Choose three.)

- A. Deploy the API to multiple Regions. Configure Amazon Route 53 with custom domain names that route traffic to each Regional API endpoint. Implement a Route 53 multivalued answer routing policy.
- B. Create a new KMS multi-Region customer managed key. Create a new KMS customer managed replica key in each in-scope Region.
- C. Replicate the existing Secrets Manager secret to other Regions. For each in-scope Region's replicated secret, select the appropriate KMS key.
- D. Create a new AWS managed KMS key in each in-scope Region. Convert an existing key to a multi-Region key. Use the multi-Region key in other Regions.
- E. Create a new Secrets Manager secret in each in-scope Region. Copy the secret value from the existing Region to the new secret in each in-scope Region.
- F. Modify the deployment process for the Lambda function to repeat the deployment across in-scope Regions. Turn on the multi-Region option for the existing API. Select the Lambda function that is deployed in each Region as the backend for the multi-Region API.

Answer: A,B,C

Explanation:

The combination of changes that will meet the requirement with the least operational overhead are: A . Deploy the API to multiple Regions. Configure Amazon Route 53 with custom domain names that route traffic to each Regional API endpoint. Implement a Route 53 multivalued answer routing policy. B . Create a new KMS multi-Region customer managed key. Create a new KMS customer managed replica key in each in-scope Region.

C . Replicate the existing Secrets Manager secret to other Regions. For each in-scope Region's replicated secret, select the appropriate KMS key.

These changes will enable the company to have an active-active configuration for its API across multiple Regions, while minimizing the complexity and cost of managing the secrets and keys.

A . This change will allow the company to use Route 53 to distribute traffic across multiple Regional API endpoints, based on the availability and latency of each endpoint. This will improve the performance and availability of the API for global customers¹²

B . This change will allow the company to use KMS multi-Region keys, which are KMS keys in different Regions that can be used interchangeably. This will simplify the encryption and decryption of secrets across Regions, as the same key material and key ID can be used in any Region³⁴

C . This change will allow the company to use Secrets Manager replication, which replicates the encrypted secret data and metadata across the specified Regions. This will ensure that the secrets are consistent and accessible in any Region, and that any update made to the primary secret will be propagated to the replica secrets automatically⁵⁶

Reference:

1:Creating a regional API endpoint - Amazon API Gateway²:Multivalue answer routing policy - Amazon Route 53³:Multi-Region keys in AWS KMS - AWS Key Management Service⁴:Creating multiRegion keys - AWS Key Management Service⁵:Replicate an AWS Secrets Manager secret to other AWS Regions⁶:How to replicate secrets in AWS Secrets Manager to multiple Regions | AWS Security Blog

Question: 358

A company has AWS accounts that are in an organization in AWS rganizations. The company wants to track Amazon EC2 usage as a metric.

The company's architecture team must receive a daily alert if the EC2 usage is more than 10% higher than the average EC2 usage from the last 30 days.

Which solution will meet these requirements?

A. Configure AWS Budgets in the organization's management account. Specify a usage type of EC2 running hours. Specify a daily period. Set the budget amount to be 10% more than the reported average usage for the last 30 days from AWS Cost Explorer.

B. Configure an alert to notify the architecture team if the usage threshold is met. Configure AWS Cost Anomaly Detection in the organization's management account. Configure a monitor type of AWS Service. Apply a filter of Amazon EC2. Configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days.

C. Enable AWS Trusted Advisor in the organization's management account. Configure a cost optimization advisory alert to notify the architecture team if the EC2 usage is 10% more than the reported average usage for the last 30 days.

D. Configure Amazon Detective in the organization's management account. Configure an EC2 usage anomaly alert to notify the architecture team if Detective identifies a usage anomaly of more than 10%.

Answer: B

Explanation:

The correct answer is B.

B . This solution meets the requirements because it uses AWS Cost Anomaly Detection, which is a feature of AWS Cost Management that uses machine learning to identify and alert on anomalous spend and usage patterns. By configuring a monitor type of AWS Service and applying a filter of Amazon EC2, the solution can track the EC2 usage as a metric across the organization's accounts. By configuring an alert subscription with a threshold of 10%, the solution can notify the architecture team via email or Amazon SNS if the EC2 usage is more than 10% higher than the average usage for the last 30 days¹²

A . This solution is incorrect because it uses AWS Budgets, which is a feature of AWS Cost Management that helps to plan and track costs and usage. However, AWS Budgets does not support usage type of EC2 running hours as a budget type. The only supported usage types are Amazon S3 storage, Amazon EC2 RI utilization, and Amazon EC2 RI coverage. Moreover, AWS Budgets does not support setting the budget amount based on the reported average usage from AWS Cost Explorer. The budget amount has to be a fixed or variable value³⁴

C . This solution is incorrect because it uses AWS Trusted Advisor, which is a feature of AWS Premium Support that provides recommendations to follow best practices for cost optimization, security, performance, and fault tolerance. However, AWS Trusted Advisor does not support configuring custom alerts based on EC2 usage or average usage for the last 30 days. The only supported alerts are based on predefined checks and thresholds that are applied to all services and resources in the account⁵⁶

D . This solution is incorrect because it uses Amazon Detective, which is a service that helps to analyze and visualize security data to investigate potential security issues. However, Amazon Detective does not support configuring EC2 usage anomaly alerts based on average usage for the last 30 days. The only supported alerts are based on GuardDuty findings and other security-related events that are detected by machine learning models⁷⁸

Reference:

1:AWS Cost Anomaly Detection - Amazon Web Services²:Getting started with AWS Cost Anomaly

Detection³:Set Custom Cost and Usage Budgets - AWS Budgets - Amazon Web Services⁴:Creating a budget - AWS Cost Management⁵:AWS Trusted Advisor⁶:AWS Trusted Advisor - AWS Support⁷:Security Investigation Visualization - Amazon Detective - AWS⁸:What is Amazon Detective? - Amazon Detective

Question: 359

An online gaming company needs to optimize the cost of its workloads on AWS. The company uses a dedicated account to host the production environment for its online gaming application and an analytics application.

Amazon EC2 instances host the gaming application and must always be available. The EC2 instances run all year. The analytics application uses data that is stored in Amazon S3. The analytics application can be interrupted and resumed without issue.

Which solution will meet these requirements MOST cost-effectively?

A. Purchase an EC2 Instance Savings Plan for the online gaming application instances. Use On- Demand

Instances for the analytics application.

- B. Purchase an EC2 Instance Savings Plan for the online gaming application instances. Use Spot Instances for the analytics application.
- C. Use Spot Instances for the online gaming application and the analytics application. Set up a catalog in AWS Service Catalog to provision services at a discount.
- D. Use On-Demand Instances for the online gaming application. Use Spot Instances for the analytics application. Set up a catalog in AWS Service Catalog to provision services at a discount.

Answer: B

Explanation:

The correct answer is B.

B . This solution is the most cost-effective because it uses an EC2 Instance Savings Plan for the online gaming application instances, which provides the lowest prices and savings up to 72% compared to On-Demand prices. The EC2 Instance Savings Plan applies to any instance size within the same family and region, regardless of availability zone, operating system, or tenancy. The online gaming application instances run all year and must always be available, so they are not suitable for Spot Instances, which can be interrupted with a two-minute notice. This solution also uses Spot Instances for the analytics application, which can reduce the cost by up to 90% compared to On-Demand prices. The analytics application can be interrupted and resumed without issue, so it is a good fit for Spot Instances, which use spare EC2 capacity. This solution does not require AWS Service Catalog, which is a service that helps to create and manage catalogs of IT services that are approved for use ON AWS, but does not provide any discounts¹²³

A . This solution is incorrect because it uses On-Demand Instances for the analytics application, which are more expensive than Spot Instances. The analytics application can be interrupted and resumed without issue, so it can benefit from the lower cost of Spot Instances, which use spare EC2 capacity.

C . This solution is incorrect because it uses Spot Instances for the online gaming application, which can be interrupted with a two-minute notice. The online gaming application instances must always be available, so they are not suitable for Spot Instances, which use spare EC2 capacity. This solution also uses AWS Service Catalog, which is a service that helps to create and manage catalogs of IT services that are approved for use on AWS, but does not provide any discounts.

D . This solution is incorrect because it uses On-Demand Instances for the online gaming application, which are more expensive than an EC2 Instance Savings Plan. The online gaming application instances run all year and must always be available, so they are suitable for an EC2 Instance Savings Plan, which provides the lowest prices and savings up to 72% compared to On-Demand prices. This solution also uses AWS Service Catalog, which is a service that helps to create and manage catalogs of IT services that are approved for use on AWS, but does not provide any discounts. Reference:

1:EC2 Instance Savings Plans - Amazon Web Services
2:Amazon EC2 Spot Instances
3:Cloud Management and Governance - AWS Service Catalog - Amazon Web Services

Question: 360

A company is preparing to deploy an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for a workload.

The company expects the cluster to support an unpredictable number of stateless pods. Many of the pods will be created during a short time period as the workload automatically scales the number of replicas that the workload uses.

Which solution will MAXIMIZE node resilience?

- A. Use a separate launch template to deploy the EKS control plane into a second cluster that is separate from the workload node groups.
- B. Update the workload node groups. Use a smaller number of node groups and larger instances in the node groups.
- C. Configure the Kubernetes Cluster Autoscaler to ensure that the compute capacity of the workload node groups stays under provisioned.
- D. Configure the workload to use topology spread constraints that are based on Availability Zone.

Answer: D

Explanation:

Configuring the workload to use topology spread constraints that are based on Availability Zone will maximize the node resilience of the workload node groups. This will ensure that the pods are evenly distributed across different Availability Zones, reducing the impact of failures or disruptions in one Availability Zone². This will also improve the availability and scalability of the workload node groups, as they can leverage the low-latency, high-throughput, and highly redundant networking between Availability Zones¹.

Question: 361

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:

Hosting the .NET Core components on Amazon ECS with the Amazon EC2 launch type will meet the requirements of having complex orchestration and full control over networking and host configuration. Amazon ECS is a fully managed container orchestration service that supports both AWS Fargate and Amazon EC2 as launch types. The Amazon EC2 launch type allows users to choose their own EC2 instances, configure their own networking settings, and access their own host operating systems. Hosting the database on Amazon Aurora MySQL Serverless v2 will meet the requirements of having a strongly relational database model and using the same database engine as SQL Server. MySQL is a compatible relational database engine with SQL Server, and it can support most of the legacy application's database model. Amazon Aurora MySQL Serverless v2 is a serverless version of Amazon Aurora MySQL that can scale up and down automatically based on demand. Using Amazon SQS for asynchronous messaging will meet the requirements of providing a compatible replacement for MSMQ, which is a queue-based messaging system³. Amazon SQS is a fully managed message queuing service that enables decoupled and scalable microservices, distributed systems, and serverless applications.

Question: 362

A company is planning to migrate its on-premises transaction-processing application to AWS. The application runs inside Docker containers that are hosted on VMS in the company's data center. The Docker containers have shared storage where the application records transaction data.

The transactions are time sensitive. The volume of transactions inside the application is unpredictable. The company must implement a low-latency storage solution that will automatically scale throughput to meet increased demand. The company cannot develop the application further and cannot continue to administer the Docker hosting environment.

How should the company migrate the application to AWS to meet these requirements?

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

instances for the EFS file system.

Answer: B

Explanation:

Migrating the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS) will meet the requirement of not administering the Docker hosting environment. AWS Fargate is a serverless compute engine that runs containers without requiring any infrastructure management³. Creating an Amazon Elastic File System (Amazon EFS) file system and adding a volume to the Fargate task definition to point to the EFS file system will meet the requirement of low-latency storage that will automatically scale throughput to meet increased demand. Amazon EFS is a fully managed file system service that provides shared access to data from multiple containers, supports NFSv4 protocol, and offers consistent performance and high availability⁴. Amazon EFS also supports automatic scaling of throughput based on the amount of data stored in the file system⁵.

Question: 363

A company is deploying a third-party web application on AWS. The application is packaged as a Docker image. The company has deployed the Docker image as an AWS Fargate service in Amazon Elastic Container Service (Amazon ECS). An Application Load Balancer (ALB) directs traffic to the application.

The company needs to give only a specific list of users the ability to access the application from the internet. The company cannot change the application and cannot integrate the application with an identity provider. All users must be authenticated through multi-factor authentication (MFA).

Which solution will meet these requirements?

- A. Create a user pool in Amazon Cognito. Configure the pool for the application. Populate the pool with the required users. Configure the pool to require MFA. Configure a listener rule on the ALB to require authentication through the Amazon Cognito hosted UI.
- B. Configure the users in AWS Identity and Access Management (IAM). Attach a resource policy to the Fargate service to require users to use MFA. Configure a listener rule on the ALB to require authentication through IAM.
- C. Configure the users in AWS Identity and Access Management (IAM). Enable AWS IAM Identity Center (AWS Single Sign-On). Configure resource protection for the ALB. Create a resource protection rule to require users to use MFA.
- D. Create a user pool in AWS Amplify. Configure the pool for the application. Populate the pool with the required users. Configure the pool to require MFA. Configure a listener rule on the ALB to require authentication through the Amplify hosted UI.

Answer: A

Explanation:

Creating a user pool in Amazon Cognito and configuring it for the application will meet the requirement of giving only a specific list of users the ability to access the application from the internet. A user pool is a directory of users that can sign in to an application with username and password¹. The company can populate the user pool with the required users and configure the pool to require MFA for additional security². Configuring a listener rule on the ALB to require authentication through the Amazon Cognito hosted UI will meet the requirement of not changing the application and not integrating it with an identity provider. The ALB can use Amazon Cognito as an authentication action to authenticate users before forwarding requests to the Fargate service³. The Amazon Cognito hosted UI is a customizable web page that provides sign-in and sign-up functionality for users⁴.

Question: 364

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

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

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

Answer: C

Explanation:

S3 Transfer Acceleration is a feature that enables fast, easy, and secure transfers of files over long distances between your client and an S3 bucket¹. It works by leveraging the CloudFront edge network to route your requests to S3 over an optimized network path¹. By using a Transfer Acceleration endpoint when generating a

presigned URL, you can allow authenticated users to upload objects faster and more reliably². Additionally, using the S3 multipart upload API can improve the performance of large object uploads by breaking them into smaller parts and uploading them in parallel³.

S3 Transfer Acceleration

Using Transfer Acceleration with presigned URLs

Uploading objects using multipart upload API

Question: 365

A company has a solution that analyzes weather data from thousands of weather stations. The weather stations send the data over an Amazon API Gateway REST API that has an AWS Lambda function integration. The Lambda function calls a third-party service for data pre-processing. The third-party service gets overloaded and fails the pre-processing, causing a loss of data.

A solutions architect must improve the resiliency of the solution. The solutions architect must ensure that no data is lost and that data can be processed later if failures occur.

What should the solutions architect do to meet these requirements?

- A. Create an Amazon Simple Queue Service (Amazon SQS) queue. Configure the queue as the dead letter queue for the API.
- B. Create two Amazon Simple Queue Service (Amazon SQS) queues: a primary queue and a secondary queue. Configure the secondary queue as the dead-letter queue for the primary queue. Update the API to use a new integration to the primary queue. Configure the Lambda function as the invocation target for the primary queue.
- C. Create two Amazon EventBridge event buses: a primary event bus and a secondary event bus. Update the API to use a new integration to the primary event bus. Configure an EventBridge rule to react to all events on the primary event bus. Specify the Lambda function as the target of the rule. Configure the secondary event bus as the failure destination for the Lambda function.
- D. Create a custom Amazon EventBridge event bus. Configure the event bus as the failure destination for the Lambda function.

Answer: C

Explanation:

This option allows the solution to decouple the API from the Lambda function and use EventBridge as an event-driven service that can handle failures gracefully¹. By using two event buses, one for normal events and one for failed events, the solution can ensure that no data is lost and that data can be processed later if failures occur². The primary event bus receives the data from the weather stations through the API integration and triggers the Lambda function through a rule. The Lambda function can then call the third-party service for data pre-processing. If the third-party service fails, the Lambda function can send an error response to EventBridge,

which will route it to the secondary event bus as a failure destination³. The secondary event bus can then store the failed events in another service, such as Amazon S3 or Amazon SQS, for troubleshooting or reprocessing.

Using Amazon EventBridge with AWS Lambda

Using multiple event buses Using failure destinations [Using dead-letter queues]

Question: 366

A company built an ecommerce website on AWS using a three-tier web architecture. The application is Java-based and composed of an Amazon CloudFront distribution, an Apache web server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database. Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts. The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics. Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis. Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events? (Choose three.)

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

Answer: A,B,D

Explanation:

Configuring the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs will allow the solutions architect to monitor and troubleshoot the database performance by identifying slow or problematic queries¹. CloudWatch Logs also provides features such as metric filters, alarms, and dashboards to analyze and visualize the log data².

Implementing the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java will allow the solutions architect to measure and map the end-to-end latency and performance of the web application³. X-Ray traces show how requests travel through the application components, such as web servers, load balancers, microservices, and databases⁴. X-Ray also provides features such as service maps, annotations, histograms, and error rates to analyze and optimize the application performance.

Installing and configuring an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs will allow the solutions architect to monitor and troubleshoot the web server performance by collecting and storing the Apache access and error logs. CloudWatch Logs also provides features such as metric

filters, alarms, and dashboards to analyze and visualize the log data².

Publishing Aurora MySQL logs to Amazon CloudWatch Logs

Working with log data in CloudWatch Logs

Instrumenting your application with the X-Ray SDK for Java

Tracing requests with AWS X-Ray [Analyzing application performance with AWS X-Ray] [Using CloudWatch Logs with your Apache web server]

Question: 367

A company provides a software as a service (SaaS) application that runs in the AWS Cloud. The application runs on Amazon EC2 instances behind a Network LoadBalancer (NLB). The instances are in an Auto Scaling group and are distributed across three Availability Zones in a single AWS Region.

The company is deploying the application into additional Regions. The company must provide static IP addresses for the application to customers so that the customers can add the IP addresses to allow lists.

The solution must automatically route customers to the Region that is geographically closest to them.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront distribution. Create a CloudFront origin group. Add the NLB for each additional Region to the origin group. Provide customers with the IP address ranges of the distribution's edge locations.
- B. Create an AWS Global Accelerator standard accelerator. Create a standard accelerator endpoint for the NLB in each additional Region. Provide customers with the Global Accelerator IP address.
- C. Create an Amazon CloudFront distribution. Create a custom origin for the NLB in each additional Region. Provide customers with the IP address ranges of the distribution's edge locations.
- D. Create an AWS Global Accelerator custom routing accelerator. Create a listener for the custom routing accelerator. Add the IP address and ports for the NLB in each additional Region. Provide customers with the Global Accelerator IP address.

Answer: B

Explanation:

AWS Global Accelerator is a networking service that helps you improve the availability and performance of the applications that you offer to your global users¹. It provides static IP addresses that act as a fixed entry point to your applications and route user traffic to the optimal endpoint based on performance, health, and policies that you configure¹. By creating a standard accelerator endpoint for the NLB in each additional Region, you can ensure that customers are automatically directed to the Region that is geographically closest to them². You can also provide customers with the Global Accelerator IP address, which is anycast from AWS edge locations and

does not change when you add or remove endpoints³.

What is AWS Global Accelerator?

Standard accelerator endpoints

AWS Global Accelerator IP addresses

Question: 368

A company is running multiple workloads in the AWS Cloud. The company has separate units for software development. The company uses AWS Organizations and federation with SAML to give permissions to developers to manage resources in their AWS accounts. The development units each deploy their production workloads into a common production account.

Recently, an incident occurred in the production account in which members of a development unit terminated an EC2 instance that belonged to a different development unit. A solutions architect must create a solution that prevents a similar incident from happening in the future. The solution also must allow developers the possibility to manage the instances used for their workloads. Which strategy will meet these requirements?

- A. Create separate OUs in AWS Organizations for each development unit. Assign the created OUs to the company AWS accounts. Create separate SCPs with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag that matches the development unit name. Assign the SCP to the corresponding OU.
- B. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation. Update the IAM policy for the developers' assumed IAM role with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag and aws:PrincipalTag/ DevelopmentUnit.
- C. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation. Create an SCP with an allow action and a StringEquals condition for the DevelopmentUnit resource tag and aws:PrincipalTag/DevelopmentUnit. Assign the SCP to the root OU.
- D. Create separate IAM policies for each development unit. For every IAM policy, add an allow action and a StringEquals condition for the DevelopmentUnit resource tag and the development unit name. During SAML federation, use AWS Security Token Service (AWS STS) to assign the IAM policy and match the development unit name to the assumed IAM role.

Answer: B

Explanation:

This option allows the solutions architect to use session tags to pass additional information about the federated user, such as the development unit name, to AWS¹. Session tags are key-value pairs that you can define in your identity provider (IdP) and pass in your SAML assertion¹. By using a deny action and a StringNotEquals condition in the IAM policy, you can prevent developers from accessing or modifying EC2 instances that belong to a different development unit². This way, you can enforce fine-grained access control and prevent accidental or malicious incidents.

Passing session tags in SAML assertions

Using tags for attribute-based access control

Question: 369

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

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

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

Provide the ability to distribute infrastructure services to multiple accounts in AWS Organizations. Provide the ability to enforce tags on any infrastructure that is started by users.

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

- A. Develop infrastructure services using AWS Cloud Formation templates. Add the templates to a central Amazon S3 bucket and add the IAM roles or users that require access to the S3 bucket policy.
- B. Develop infrastructure services using AWS Cloud Formation templates. Upload each template as an AWS Service Catalog product to portfolios created in a central AWS account. Share these portfolios with the Organizations structure created for the company.
- C. Allow user IAM roles to have AWSCloudFormationFullAccess and AmazonS3ReadOnlyAccess permissions. Add an Organizations SCP at the AWS account root user level to deny all services except AWS CloudFormation and Amazon S3.
- D. Allow user IAM roles to have ServiceCatalogEndUserAccess permissions only. Use an automation script to import the central portfolios to local AWS accounts, copy the TagOption, assign users access, and apply launch constraints.
- E. Use the AWS Service Catalog TagOption Library to maintain a list of tags required by the company. Apply the TagOption to AWS Service Catalog products or portfolios.
- F. Use the AWS CloudFormation Resource Tags property to enforce the application of tags to any CloudFormation templates that will be created for users.

Answer: B,D,E

Explanation:

Developing infrastructure services using AWS CloudFormation templates and uploading them as AWS Service Catalog products to portfolios created in a central AWS account will enable the company to centrally manage the creation of infrastructure services and control who can use them¹. AWS Service Catalog allows you to create and manage catalogs of IT services that are approved for use on AWS². You can organize products into portfolios, which are collections of products along with configuration information³. You can share portfolios with other accounts in your organization using AWS Organizations⁴.

Allowing user IAM roles to have ServiceCatalogEndUserAccess permissions only and using an automation script to import the central portfolios to local AWS accounts, copy the TagOption, assign users access, and apply launch constraints will enable the company to provide least privilege access to users when launching AWS infrastructure services. ServiceCatalogEndUserAccess is a managed IAM policy that grants users permission to list and view products and launch product instances. An automation script can help import the shared portfolios from the central account to the local accounts, copy the TagOption from the central account, assign users access to the portfolios, and apply launch constraints that specify which IAM role or user can provision a

product.

Using the AWS Service Catalog TagOption Library to maintain a list of tags required by the company and applying the TagOption to AWS Service Catalog products or portfolios will enable the company to enforce tags on any infrastructure that is started by users. TagOptions are key-value pairs that you can use to classify your AWS Service Catalog resources. You can create a TagOption Library that contains all the tags that you want to use across your organization. You can apply TagOptions to products or portfolios, and they will be automatically applied to any provisioned product instances.

Creating a product from an existing CloudFormation template

What is AWS Service Catalog?

Working with portfolios

Sharing a portfolio with AWS Organizations [Providing least privilege access for users] [AWS managed policies for job functions] [Importing shared portfolios] [Enforcing tag policies] [Working with TagOptions] [Creating a TagOption Library] [Applying TagOptions]

Question: 370

A car rental company has built a serverless REST API to provide data to its mobile app. The app consists of an Amazon API Gateway API with a Regional endpoint, AWS Lambda functions, and an Amazon Aurora MySQL Serverless DB cluster. The company recently opened the API to mobile apps of partners. A significant increase in the number of requests resulted, causing sporadic database memory errors. Analysis of the API traffic indicates that clients are making multiple HTTP GET requests for the same queries in a short period of time. Traffic is concentrated during business hours, with spikes around holidays and other events.

The company needs to improve its ability to support the additional usage while minimizing the increase in costs associated with the solution.

Which strategy meets these requirements?

- A. Convert the API Gateway Regional endpoint to an edge-optimized endpoint. Enable caching in the production stage.
- B. Implement an Amazon ElastiCache for Redis cache to store the results of the database calls. Modify the Lambda functions to use the cache.
- C. Modify the Aurora Serverless DB cluster configuration to increase the maximum amount of available memory.
- D. Enable throttling in the API Gateway production stage. Set the rate and burst values to limit the incoming calls.

Answer: A

Explanation:

This option allows the company to use Amazon CloudFront to improve the latency and availability of the API requests by caching the responses at the edge locations closest to the clients¹. By enabling caching in the production stage, the company can reduce the number of calls made to the backend services, such as Lambda functions and Aurora Serverless DB cluster, and save on costs and resources². This option also helps to handle

traffic spikes and reduce database memory errors by serving cached responses instead of querying the database repeatedly.

Choosing an API endpoint type

Enabling API caching to enhance responsiveness

Question: 371

A company hosts a software as a service (SaaS) solution on AWS. The solution has an Amazon API Gateway API that serves an HTTPS endpoint. The API uses AWS Lambda functions for compute. The Lambda functions store data in an Amazon Aurora Serverless V1 database.

The company used the AWS Serverless Application Model (AWS SAM) to deploy the solution. The solution extends across multiple Availability Zones and has a disaster recovery (DR) plan.

A solutions architect must design a DR strategy that can recover the solution in another AWS Region. The solution has an RTO of 5 minutes and an RPO of 1 minute.

What should the solutions architect do to meet these requirements?

- A. Create a read replica of the Aurora Serverless V1 database in the target Region. Use AWS SAM to create a runbook to deploy the solution to the target Region. Promote the read replica to primary in case of disaster.
- B. Change the Aurora Serverless V1 database to a standard Aurora MySQL global database that extends across the source Region and the target Region. Use AWS SAM to create a runbook to deploy the solution to the target Region.
- C. Create an Aurora Serverless V1 DB cluster that has multiple writer instances in the target Region. Launch the solution in the target Region. Configure the two Regional solutions to work in an active-passive configuration.
- D. Change the Aurora Serverless V1 database to a standard Aurora MySQL global database that extends across the source Region and the target Region. Launch the solution in the target Region. Configure the two Regional solutions to work in an active-passive configuration.

Answer: D

Explanation:

This option allows the solutions architect to use Aurora global database to replicate data across multiple AWS Regions with low latency and high availability¹. By launching the solution in the target Region, the solutions architect can ensure that the API Gateway, Lambda functions, and other resources are ready to serve traffic in case of a disaster in the source Region. By configuring the two Regional solutions to work in an active-passive configuration, the solutions architect can minimize costs and avoid data conflicts by having only one primary Region that accepts write operations and one secondary Region that serves as a standby². The RTO and RPO requirements can be met by using Aurora global database, which supports sub-second failover times and near real-time replication¹.

Question: 372

A company owns a chain of travel agencies and is running an application in the AWS Cloud. Company employees use the application to search for information about travel destinations. Destination content is updated four times each year. Two fixed Amazon EC2 instances serve the application. The company uses an Amazon Route 53 public hosted zone with a multivalued record of travel.example.com that returns the Elastic IP addresses for the EC2 instances. The application uses Amazon DynamoDB as its primary data store. The company uses a self-hosted Redis instance as a caching solution.

During content updates, the load on the EC2 instances and the caching solution increases drastically. This increased load has led to downtime on several occasions. A solutions architect must update the application so that the application is highly available and can handle the load that is generated by the content updates. Which solution will meet these requirements?

- A. Set up DynamoDB Accelerator (DAX) as in-memory cache. Update the application to use DAX. Create an Auto Scaling group for the EC2 instances. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the ALB. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alias. Configure scheduled scaling for the EC2 instances before the content updates.
- B. Set up Amazon ElastiCache for Redis. Update the application to use ElastiCache. Create an Auto Scaling group for the EC2 instances. Create an Amazon CloudFront distribution, and set the Auto Scaling group as an origin for the distribution. Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alias. Manually scale up EC2 instances before the content updates.
- C. Set up Amazon ElastiCache for Memcached. Update the application to use ElastiCache. Create an Auto Scaling group for the EC2 instances. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the ALB. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alias. Configure scheduled scaling for the application before the content updates.
- D. Set up DynamoDB Accelerator (DAX) as in-memory cache. Update the application to use DAX. Create an Auto Scaling group for the EC2 instances. Create an Amazon CloudFront distribution, and set the Auto Scaling group as an origin for the distribution. Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alias. Manually scale up EC2 instances before the content updates.

Answer: A

Explanation:

This option allows the company to use DAX to improve the performance and reduce the latency of the DynamoDB queries by caching the results in memory¹. By updating the application to use DAX, the company can reduce the load on the DynamoDB tables and avoid throttling errors¹. By creating an Auto Scaling group for the EC2 instances, the company can adjust the number of instances based on the demand and ensure high availability². By creating an ALB, the company can distribute the incoming traffic across multiple EC2 instances and improve fault tolerance³. By updating the Route 53 record to use a simple routing policy that targets the ALB's DNS alias, the company can route users to the ALB endpoint and leverage its health checks and load

balancing features⁴. By configuring scheduled scaling for the EC2 instances before the content updates, the company can anticipate and handle traffic spikes during peak periods⁵.

What is Amazon DynamoDB Accelerator (DAX)?

What is Amazon EC2 Auto Scaling?

What is an Application Load Balancer?

Choosing a routing policy

Scheduled scaling for Amazon EC2 Auto Scaling

Question: 373

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? (Select 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: B,C,E

Explanation:

Resource-based policies are policies that you attach to a resource, such as an IAM role, to specify who can access the resource and what actions they can perform on it¹. Identity-based policies are policies that you attach to an IAM user, group, or role to specify what actions they can perform on which resources². Trust policies are special types of resource-based policies that define which principals (such as IAM users or roles) can assume a role³.

To allow an IAM user in Account A to assume a role in Account B, the solutions architect needs to do the following:

Configure the resource-based policy on the target role in Account B to allow the action `sts:AssumeRole` for the IAM user in Account A. This policy grants permission to the IAM user to assume the role⁴.

Configure the identity-based policy on the user in Account A to allow the action `sts:AssumeRole` for the target role in Account B. This policy grants permission to the user to perform the action of assuming the role⁵.

Configure the trust policy on the target role in Account B to allow the principal of the IAM user in Account A. This policy defines who can assume the role.

Resource-based policies

Identity-based policies

Trust policies

Granting a user permissions to switch roles

Switching roles

[\[Modifying a role trust policy\]](#)

Question: 374

A company uses AWS Organizations to manage a multi-account structure. The company has hundreds of AWS accounts and expects the number of accounts to increase. The company is building a new application that uses Docker images. The company will push the Docker images to Amazon Elastic Container Registry (Amazon ECR). Only accounts that are within the company's organization should have access to the images.

The company has a CI/CD process that runs frequently. The company wants to retain all the tagged images. However, the company wants to retain only the five most recent untagged images.

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

- A. Create a private repository in Amazon ECR. Create a permissions policy for the repository that allows only required ECR operations. Include a condition to allow the ECR operations if the value of the `aws:PrincipalOrgID` condition key is equal to the ID of the company's organization. Add a lifecycle rule to the ECR repository that deletes all untagged images over the count of five.
- B. Create a public repository in Amazon ECR. Create an IAM role in the ECR account. Set permissions so that any account can assume the role if the value of the `aws:PrincipalOrgID` condition key is equal to the ID of the company's organization. Add a lifecycle rule to the ECR repository that deletes all untagged images over the count of five.
- C. Create a private repository in Amazon ECR. Create a permissions policy for the repository that includes only required ECR operations. Include a condition to allow the ECR operations for all account IDs in the organization. Schedule a daily Amazon EventBridge rule to invoke an AWS Lambda function that deletes all untagged images over the count of five.
- D. Create a public repository in Amazon ECR. Configure Amazon ECR to use an interface VPC endpoint with an endpoint policy that includes the required permissions for images that the company needs to pull. Include a condition to allow the ECR operations for all account IDs in the company's organization. Schedule a daily Amazon EventBridge rule to invoke an AWS Lambda function that deletes all untagged images over the count of five.

Answer: A

Explanation:

This option allows the company to use a private repository in Amazon ECR to store and manage its Docker images securely and efficiently¹. By creating a permissions policy for the repository that allows only required ECR operations, such as `ecr:GetDownloadUrlForLayer`, `ecr:BatchGetImage`, `ecr:BatchCheckLayerAvailability`, `ecr:PutImage`, and `ecr:InitiateLayerUpload`², the company can restrict access to the repository and prevent unauthorized actions. By including a condition to allow the ECR operations if the value of the `aws:PrincipalOrgID` condition key is equal to the ID of the company's organization, the company can ensure that only accounts that are within its organization can access the images³. By adding a lifecycle rule to the ECR repository that deletes all untagged images over the count of five, the company can reduce storage costs and retain only the most recent untagged images⁴.

Amazon ECR private repositories

Amazon ECR repository policies

Restricting access to AWS Organizations members

Amazon ECR lifecycle policies

Question: 375

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

Scaling group.

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

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

- A. Create an AWS CloudFormation template that contains a transit gateway attachment and related routing configurations. Create a CloudFormation stack set that includes this template. Use CloudFormation StackSets to deploy a new stack for each VPC in the account. Deploy a new VPC for each test environment.
- B. Create a single VPC for the test environments. Include a transit gateway attachment and related routing configurations. Use AWS CloudFormation to deploy all test environments into the VPC.
- C. Create a new OU in AWS Organizations for testing. Create an AWS CloudFormation template that contains a VPC, necessary networking resources, a transit gateway attachment, and related routing configurations. Create a CloudFormation stack set that includes this template. Use CloudFormation StackSets for deployments into each account under the testing OU. Create a new account for each test environment.
- D. Convert the test environment EC2 instances into Docker images. Use AWS CloudFormation to configure an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in a new VPC, create a transit gateway attachment, and create related routing configurations. Use Kubernetes to manage the deployment and lifecycle of the test environments.

Answer: B

Explanation:

This option allows the company to use a single VPC to host multiple test environments that are isolated from each other by using different subnets and security groups¹. By including a transit gateway attachment and related routing configurations, the company can enable the test environments to communicate with the central server in the on-premises data center through a VPN connection². By using AWS CloudFormation to deploy all test environments into the VPC, the company can automate the creation and deletion of test environments without any manual intervention³. This option also minimizes the operational overhead by reducing the number of VPCs, accounts, and resources that need to be managed.

Working with VPCs and subnets Working with transit gateways Working with AWS CloudFormation stacks

Question: 376

A solutions architect is determining the DNS strategy for an existing VPC. The VPC is provisioned to use the 10.24.34.0/24 CIDR block. The VPC also uses Amazon Route 53 Resolver for DNS. New requirements mandate that DNS queries must use private hosted zones. Additionally, instances that have public IP addresses must receive corresponding public hostnames.

Which solution will meet these requirements to ensure that the domain names are correctly resolved within the VPC?

- A. Create a private hosted zone. Activate the enableDnsSupport attribute and the enableDnsHostnames attribute for the VPC. Update the VPC DHCP options set to include domain-name-servers=10.24.34.2.
- B. Create a private hosted zone. Associate the private hosted zone with the VPC. Activate the enableDnsSupport attribute and the enableDnsHostnames attribute for the VPC. Create a new VPC DHCP options set, and configure domain-name-servers=AmazonProvidedDNS. Associate the new DHCP options set with the VPC.
- C. Deactivate the enableDnsSupport attribute for the VPC. Activate the enableDnsHostnames attribute for the VPC. Create a new VPC DHCP options set, and configure domain-name-servers=10.24.34.2. Associate the new DHCP options set with the VPC.
- D. Create a private hosted zone. Associate the private hosted zone with the VPC. Activate the enableDnsSupport attribute for the VPC. Deactivate the enableDnsHostnames attribute for the VPC. Update the VPC DHCP options set to include domain-name-servers=AmazonProvidedDNS.

Answer: B

Explanation:

This option allows the solutions architect to use a private hosted zone to host DNS records that are only accessible within the VPC¹. By associating the private hosted zone with the VPC, the solutions architect can ensure that DNS queries from the VPC are routed to the private hosted zone². By activating the enableDnsSupport attribute and the enableDnsHostnames attribute for the VPC, the solutions architect can

enable DNS resolution and hostname assignment for instances in the VPC3. By creating a new VPC DHCP options set, and configuring domain-name-servers=AmazonProvidedDNS, the solutions architect can use Amazon-provided DNS servers to resolve DNS queries from instances in the VPC4. By associating the new DHCP options set with the VPC, the solutions architect can apply the DNS settings to all instances in the VPC5.

What is Amazon Route 53 Resolver?

Associating a private hosted zone with your VPC

Using DNS with your VPC

DHCP options sets

Modifying your DHCP options

Question: 377

A financial services company runs a complex, multi-tier application on Amazon EC2 instances and AWS Lambda functions. The application stores temporary data in Amazon S3. The S3 objects are valid for only 45 minutes and are deleted after 24 hours.

The company deploys each version of the application by launching an AWS CloudFormation stack. The stack creates all resources that are required to run the application. When the company deploys and validates a new application version, the company deletes the CloudFormation stack of the old version.

The company recently tried to delete the CloudFormation stack of an old application version, but the operation failed. An analysis shows that CloudFormation failed to delete an existing S3 bucket. A solutions architect needs to resolve this issue without making major changes to the application's architecture.

Which solution meets these requirements?

- A. Implement a Lambda function that deletes all files from a given S3 bucket. Integrate this Lambda function as a custom resource into the CloudFormation stack. Ensure that the custom resource has a **DependsOn** attribute that points to the S3 bucket's resource.
- B. Modify the CloudFormation template to provision an Amazon Elastic File System (Amazon EFS) file system to store the temporary files there instead of in Amazon S3. Configure the Lambda functions to run in the same VPC as the file system. Mount the file system to the EC2 instances and Lambda functions.
- C. Modify the CloudFormation stack to create an S3 Lifecycle rule that expires all objects 45 minutes after creation. Add a **DependsOn** attribute that points to the S3 bucket's resource.
- D. Modify the CloudFormation stack to attach a **DeletionPolicy** attribute with a value of **Delete** to the S3 bucket.

Answer: D

Explanation:

This option allows the solutions architect to use a DeletionPolicy attribute to specify how AWS CloudFormation handles the deletion of an S3 bucket when the stack is deleted¹. By setting the value of Delete, the solutions architect can instruct CloudFormation to delete the bucket and all of its contents¹. This option does not require any major changes to the application's architecture or any additional resources.

Deletion policies

Question: 378

An online retail company hosts its stateful web-based application and MySQL database in an on-premises data center on a single server. The company wants to increase its customer base by conducting more marketing campaigns and promotions. In preparation, the company wants to migrate its application and database to AWS to increase the reliability of its architecture.

Which solution should provide the HIGHEST level of reliability?

- A. Migrate the database to an Amazon RDS MySQL Multi-AZ DB instance. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balancer. Store sessions in Amazon Neptune.
- B. Migrate the database to Amazon Aurora MySQL. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balancer. Store sessions in an Amazon ElastiCache for Redis replication group.
- C. Migrate the database to Amazon DocumentDB (with MongoDB compatibility). Deploy the application in an Auto Scaling group on Amazon EC2 instances behind a Network Load Balancer. Store sessions in Amazon Kinesis Data Firehose.
- D. Migrate the database to an Amazon RDS MariaDB Multi-AZ DB instance. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balancer. Store sessions in Amazon ElastiCache for Memcached.

Answer: B

Explanation:

This option allows the company to use Amazon Aurora MySQL, which is a fully managed relational database service that is compatible with MySQL and offers up to five times better performance than standard MySQL¹. By migrating the database to Aurora MySQL, the company can benefit from its high availability, durability, scalability, and security features¹. By deploying the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balancer, the company can ensure that the application can handle varying levels of traffic and distribute the requests across multiple instances². By storing sessions in an Amazon ElastiCache for Redis replication group, the company can improve the performance and reliability of the session data by using a fast, in-memory data store that supports replication and failover³.

What is Amazon Aurora? What is Auto Scaling?

What is Amazon ElastiCache?

Question: 379

A company is building an application on AWS. The application sends logs to an Amazon Elasticsearch Service (Amazon ES) cluster for analysis. All data must be stored within a VPC.

Some of the company's developers work from home. Other developers work from three different company office locations. The developers need to access Amazon ES to analyze and visualize logs directly from their local development machines.

Which solution will meet these requirements?

- A. Configure and set up an AWS Client VPN endpoint. Associate the Client VPN endpoint with a subnet in the VPC. Configure a Client VPN self-service portal. Instruct the developers to connect by using the client for Client VPN.
- B. Create a transit gateway, and connect it to the VPC. Create an AWS Site-to-Site VPN. Create an attachment to the transit gateway. Instruct the developers to connect by using an OpenVPN client.
- C. Create a transit gateway, and connect it to the VPC. Order an AWS Direct Connect connection. Set up a public VIF on the Direct Connect connection. Associate the public VIF with the transit gateway. Instruct the developers to connect to the Direct Connect connection.
- D. Create and configure a bastion host in a public subnet of the VPC. Configure the bastion host security group to allow SSH access from the company CIDR ranges. Instruct the developers to connect by using SSH.

Answer: A**Explanation:**

This option allows the company to use AWS Client VPN to enable secure and private access to the Amazon ES cluster from any location¹. By configuring and setting up an AWS Client VPN endpoint, the company can create a secure tunnel between the developers' devices and the VPC². By associating the Client VPN endpoint with a subnet in the VPC, the company can ensure that the traffic from the developers' devices is routed to the Amazon ES cluster within the VPC³. By configuring a Client VPN self-service portal, the company can enable the developers to download and install the client for Client VPN, which is based on OpenVPN⁴. By instructing the developers to connect by using the client for Client VPN, the company can allow them to access Amazon ES to analyze and visualize logs directly from their local development machines.

What is AWS Client VPN?

Creating a Client VPN endpoint

Associating a target network with a Client VPN endpoint

Configuring a self-service portal

Question: 380

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? (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:

This option allows the solutions architect to use an identity policy that grants the user read and write access to their own personal folder on the S3 bucket¹. By adding a condition that specifies that the S3 paths must be prefixed with `{aws:username}`, the solutions architect can ensure that each scientist can only access their own folder². By applying the policy on the scientists' IAM user group, the solutions architect can simplify the management of permissions for all the scientists³. By configuring a trail with AWS CloudTrail to capture all object-level events in the S3 bucket, the solutions architect can record and store information about every action performed on the S3 objects⁴. By storing the trail output in another S3 bucket, the solutions architect can secure and archive the log files⁵. By using Amazon Athena to query the logs and generate reports, the solutions architect can use a serverless interactive query service that can analyze data in S3 using standard SQL.

Identity-based policies

Policy variables

IAM groups

Object-level logging

Creating a trail that applies to all regions

[\[What is Amazon Athena?\]](#)

Question: 381

A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured

according to best practice, using least privilege.

A Solutions Architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster.

What steps are required after the deployment to meet the requirements? (Choose two.)

- A. Create tasks using the bridge network mode.
- B. Create tasks using the awsvpc network mode.
- C. Apply security groups to Amazon EC2 instances, and use IAM roles for EC2 instances to access other resources.
- D. Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources.
- E. Apply security groups to the tasks, and use IAM roles for tasks to access other resources.

Answer: B,E

Explanation:

The awsvpc network mode provides each task with its own elastic network interface (ENI) and a primary private IP address¹. By using this network mode, the solutions architect can isolate the tasks from each other and apply security groups to the tasks directly². This way, the solutions architect can control the inbound and outbound traffic at the task level and enforce the least privilege principle³. IAM roles for tasks allow the solutions architect to assign permissions to each task separately, so that they can access other AWS resources that they need⁴. By using IAM roles for tasks, the solutions architect can avoid passing IAM credentials into the container at launch time, which is less secure and more prone to errors⁵.

awsvpc network mode

Task networking with the awsvpc network mode

Security groups for your VPC

IAM roles for tasks

Best practices for managing AWS access keys

Question: 382

A company is running a serverless application that consists of several AWS Lambda functions and Amazon DynamoDB tables. The company has created new functionality that requires the Lambda functions to access an Amazon Neptune DB cluster. The Neptune DB cluster is located in three subnets in a VPC.

Which of the possible solutions will allow the Lambda functions to access the Neptune DB cluster and DynamoDB tables? (Select TWO.)

- A. Create three public subnets in the Neptune VPC, and route traffic through an internet gateway. Host the Lambda functions in the three new public subnets.
- B. Create three private subnets in the Neptune VPC, and route internet traffic through a NAT gateway. Host the Lambda functions in the three new private subnets.
- C. Host the Lambda functions outside the VPC. Update the Neptune security group to allow access from the IP ranges of the Lambda functions.

- D. Host the Lambda functions outside the VPC. Create a VPC endpoint for the Neptune database, and have the Lambda functions access Neptune over the VPC endpoint.
- E. Create three private subnets in the Neptune VPC. Host the Lambda functions in the three new isolated subnets. Create a VPC endpoint for DynamoDB, and route DynamoDB traffic to the VPC endpoint.

Answer: B,E

Explanation:

This option allows the company to use private subnets and VPC endpoints to connect the Lambda functions to the Neptune DB cluster and DynamoDB tables securely and efficiently¹. By creating three private subnets in the Neptune VPC, the company can isolate the Lambda functions from the public internet and reduce the attack surface². By routing internet traffic through a NAT gateway, the company can enable the Lambda functions to access AWS services that are outside the VPC, such as Amazon S3 or Amazon CloudWatch³. By hosting the Lambda functions in the three new private subnets, the company can ensure that the Lambda functions can access the Neptune DB cluster within the same VPC⁴. By creating a VPC endpoint for DynamoDB, the company can enable the Lambda functions to access DynamoDB tables without going through the internet or a NAT gateway⁵. By routing DynamoDB traffic to the VPC endpoint, the company can improve the performance and availability of the DynamoDB access⁵.

Configuring a Lambda function to access resources in a VPC

Working with VPCs and subnets

NAT gateways

Accessing Amazon Neptune from AWS Lambda

VPC endpoints for DynamoDB

Question: 383 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 meet the company's needs with the LEAST effort? (Select 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:

Option B is correct because AWS Organizations allows a company to create a new organization from a chosen payer account and define an organizational unit hierarchy. This way, the finance department can have a centralized method for payment but also maintain visibility into each group's spending to allocate costs. The company can also invite the existing accounts to join the organization and create new accounts using Organizations, which simplifies the account management process.

Option D is correct because enabling all features of AWS Organizations and establishing appropriate service control policies (SCPs) that filter IAM permissions for sub-accounts allows the security team to have a centralized mechanism to control IAM usage in all the company's accounts. SCPs are policies that specify the maximum permissions for an organization or organizational unit (OU), and they can be used to restrict access to certain services or actions across all accounts in an organization.

Option A is incorrect because using a collection of parameterized AWS CloudFormation templates defining common IAM permissions that are launched into each account requires more effort than using SCPs. Moreover, it does not provide a centralized mechanism to control IAM usage, as each account would have to launch the appropriate stacks to enforce the least privilege model.

Option C is incorrect because requiring each business unit to use its own AWS accounts does not provide a centralized method for payment or a centralized mechanism to control IAM usage. Tagging each AWS account appropriately and enabling Cost Explorer to administer chargebacks may help with cost allocation, but it is not as efficient as using AWS Organizations.

Option E is incorrect because consolidating all of the company's AWS accounts into a single AWS account does not provide visibility into each group's spending or a way to control IAM usage for different business units. Using tags for billing purposes and the IAM's Access Advisor feature to enforce the least privilege model may help with cost optimization and security, but it is not as scalable or flexible as using AWS Organizations.

AWS Organizations

Service Control Policies

AWS CloudFormation

Cost Explorer

IAM Access Advisor

Question: 384

A company that provisions job boards for a seasonal workforce is seeing an increase in traffic and usage. The backend services run on a pair of Amazon EC2 instances behind an Application Load Balancer with Amazon DynamoDB as the datastore. Application read and write traffic is slow during peak seasons.

Which option provides a scalable application architecture to handle peak seasons with the LEAST development effort?

- A. Migrate the backend services to AWS Lambda. Increase the read and write capacity of DynamoDB.
- B. Migrate the backend services to AWS Lambda. Configure DynamoDB to use global tables.
- C. Use Auto Scaling groups for the backend services. Use DynamoDB auto scaling.
- D. Use Auto Scaling groups for the backend services. Use Amazon Simple Queue Service (Amazon SQS) and an AWS Lambda function to write to DynamoDB.

Answer: C

Explanation:

Option C is correct because using Auto Scaling groups for the backend services allows the company to scale up or down the number of EC2 instances based on the demand and traffic. This way, the backend services can handle more requests during peak seasons without compromising performance or availability. Using DynamoDB auto scaling allows the company to adjust the provisioned read and write capacity of the table or index automatically based on the actual traffic patterns. This way, the table or index can handle sudden increases or decreases in workload without throttling or overprovisioning¹.

Option A is incorrect because migrating the backend services to AWS Lambda may require significant development effort to rewrite the code and test the functionality. Moreover, increasing the read and write capacity of DynamoDB manually may not be efficient or cost-effective, as it does not account for the variability of the workload. The company may end up paying for unused capacity or experiencing throttling if the workload exceeds the provisioned capacity¹.

Option B is incorrect because migrating the backend services to AWS Lambda may require significant development effort to rewrite the code and test the functionality. Moreover, configuring DynamoDB to use global tables may not be necessary or beneficial for the company, as global tables are mainly used for replicating data across multiple AWS Regions for fast local access and disaster recovery. Global tables do not automatically scale the provisioned capacity of each replica table; they still require manual or auto scaling settings².

Option D is incorrect because using Amazon Simple Queue Service (Amazon SQS) and an AWS Lambda function to write to DynamoDB may introduce additional complexity and latency to the application architecture. Amazon SQS is a message queue service that decouples and coordinates the components of a distributed system. AWS Lambda is a serverless compute service that runs code in response to events. Using these services may require significant development effort to integrate them with the backend services and DynamoDB. Moreover, they may not improve the read performance of DynamoDB, which may also be affected by high traffic³.

[Auto Scaling groups](#)

[DynamoDB auto scaling](#)

[AWS Lambda](#)

[DynamoDB global tables](#)

[AWS Lambda vs EC2: Comparison of AWS Compute Resources - Simform](#)

[Managing throughput capacity automatically with DynamoDB auto scaling - Amazon DynamoDB](#)

[AWS Aurora Global Database vs. DynamoDB Global Tables](#)

[Amazon Simple Queue Service \(SQS\)](#)

Question: 385

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:

The AWS Application Discovery Service can help plan migration projects by collecting data about on-premises servers, such as configuration, performance, and network connections. The data collection agent is a lightweight software that can be installed on each server to gather this information. This option is more cost-effective than agentless discovery, which requires deploying a virtual appliance in the VMware environment, or using CloudWatch agent, which incurs additional charges for CloudWatch Logs. Scanning the servers over a VPN is not a valid option for AWS Application Discovery Service. Reference: What is AWS Application Discovery Service?, Data collection methods

Question: 386

A company has a website that runs on four Amazon EC2 instances that are behind an Application Load Balancer (ALB). When the ALB detects that an EC2 instance is no longer available, an Amazon CloudWatch alarm enters the ALARM state. A member of the company's operations team then manually adds a new EC2 instance behind the ALB.

A solutions architect needs to design a highly available solution that automatically handles the replacement of EC2 instances. The company needs to minimize downtime during the switch to the new solution.

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

- A. Delete the existing ALB. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Create a new ALB. Attach the Auto Scaling group to the new ALB. Attach the existing EC2 instances to the Auto Scaling group.
- B. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Attach the Auto Scaling group to the existing ALB. Attach the existing EC2 instances to the Auto Scaling group.
- C. Delete the existing ALB and the EC2 instances. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Create a new ALB. Attach the Auto Scaling group to the new ALB. Wait for the Auto Scaling group to launch the minimum number of EC2 instances.
- D. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Attach the Auto Scaling group to the existing ALB. Wait for the existing ALB to register the existing EC2 instances with the Auto Scaling group.

Answer: B

Explanation:

The Auto Scaling group can automatically launch and terminate EC2 instances based on the demand and health of the web application. The launch template can specify the configuration of the EC2 instances, such as the AMI, instance type, security group, and user data. The existing ALB can distribute the traffic to the EC2 instances in the Auto Scaling group. The existing EC2 instances can be attached to the Auto Scaling group without deleting them or the ALB. This option minimizes downtime and preserves the current setup of the web application.

Reference: [What is Amazon EC2 Auto Scaling?], [Launch templates], [Attach a load balancer to your Auto Scaling group], [Attach EC2 instances to your Auto Scaling group]

Question: 387

A company deploys a new web application. As part of the setup, the company configures AWS WAF to log to Amazon S3 through Amazon Kinesis Data Firehose. The company develops an Amazon Athena query that runs once daily to return AWS WAF log data from the previous 24 hours. The volume of daily logs is constant.

However, over time, the same query is taking more time to run.

A solutions architect needs to design a solution to prevent the query time from continuing to increase. The solution must minimize operational overhead.

Which solution will meet these requirements?

- A. Create an AWS Lambda function that consolidates each day's AWS WAF logs into one log file.
- B. Reduce the amount of data scanned by configuring AWS WAF to send logs to a different S3 bucket each day.
- C. Update the Kinesis Data Firehose configuration to partition the data in Amazon S3 by date and time. Create external tables for Amazon Redshift. Configure Amazon Redshift Spectrum to query the data source.
- D. Modify the Kinesis Data Firehose configuration and Athena table definition to partition the data by date and time. Change the Athena query to view the relevant partitions.

Answer: D

Explanation:

The best solution is to modify the Kinesis Data Firehose configuration and Athena table definition to partition the data by date and time. This will reduce the amount of data scanned by Athena and improve the query performance. Changing the Athena query to view the relevant partitions will also help to filter out unnecessary data. This solution requires minimal operational overhead as it does not involve creating additional resources or changing the log format. Reference: [AWS WAF Developer Guide], [Amazon Kinesis Data Firehose User Guide], [Amazon Athena User Guide]

Question: 388

A company is developing a web application that runs on Amazon EC2 instances in an Auto Scaling group behind a public-facing Application Load Balancer (ALB). Only users from a specific country are allowed to access the

application. The company needs the ability to log the access requests that have been blocked. The solution should require the least possible maintenance.

Which solution meets these requirements?

- A. Create an IPSet containing a list of IP ranges that belong to the specified country. Create an AWS WAF web ACL. Configure a rule to block any requests that do not originate from an IP range in the IPSet. Associate the rule with the web ACL. Associate the web ACL with the ALB.
- B. Create an AWS WAF web ACL. Configure a rule to block any requests that do not originate from the specified country. Associate the rule with the web ACL. Associate the web ACL with the ALB.
- C. Configure AWS Shield to block any requests that do not originate from the specified country. Associate AWS Shield with the ALB.
- D. Create a security group rule that allows ports 80 and 443 from IP ranges that belong to the specified country. Associate the security group with the ALB.

Answer: B

Explanation:

The best solution is to create an AWS WAF web ACL and configure a rule to block any requests that do not originate from the specified country. This will ensure that only users from the allowed country can access the application. AWS WAF also provides logging capabilities that can capture the access requests that have been blocked. This solution requires the least possible maintenance as it does not involve updating IP ranges or security group rules. Reference: [AWS WAF Developer Guide], [AWS Shield Developer Guide]

Question: 389

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 in 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:

The best solution is to create an FSx for Windows File Server file system in the DR Region and establish connectivity between the VPCs in both Regions by using VPC peering. This will ensure that the data does not travel across the public internet and meets the compliance requirements. By using AWS DataSync with interface VPC endpoints and AWS PrivateLink, the data can be copied securely and efficiently between the FSx for Windows File Server file systems in both Regions. This solution also provides the ability to fail over to the DR Region in case of a disaster. Reference: [Amazon FSx for Windows File Server User Guide], [AWS DataSync User Guide], [Amazon VPC User Guide]

Question: 390

A company is currently in the design phase of an application that will need an RPO of less than 5 minutes and an RTO of less than 10 minutes. The solutions architecture team is forecasting that the database will store approximately 10 TB of data.

a. As part of the design, they are looking for a database solution that will provide the company with the ability to fail over to a secondary Region.

Which solution will meet these business requirements at the LOWEST cost?

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

Answer: B

Explanation:

The best solution is to deploy an Amazon RDS instance with a cross-Region read replica in a secondary Region. This will provide the company with a database solution that can fail over to the secondary Region in case of a disaster. The read replica will have minimal replication lag and can be promoted to become the primary in less than 10 minutes, meeting the RTO requirement. The RPO requirement of less than 5 minutes can also be met by using synchronous replication within the primary Region and asynchronous replication across Regions. This solution will also have the lowest cost compared to the other options, as it does not involve

additional services or resources. Reference: [Amazon RDS User Guide], [Amazon Aurora User Guide]

Question: 391

A financial company needs to create a separate AWS account for a new digital wallet application. The company uses AWS Organizations to manage its accounts. A solutions architect uses the IAM user Support1 from the management account to create a new member account with `finance1@example.com` as the email address.

What should the solutions architect do to create IAM users in the new member account?

- A. Sign in to the AWS Management Console with AWS account root user credentials by using the 64-character password from the initial AWS Organizations email `senttofinance1@example.com`. Set up the IAM users as required.
- B. From the management account, switch roles to assume the `OrganizationAccountAccessRole` role with the account ID of the new member account. Set up the IAM users as required.
- C. Go to the AWS Management Console sign-in page. Choose "Sign in using root account credentials." Sign in by using the email address `finance1@example.com` and the management account's root password. Set up the IAM users as required.
- D. Go to the AWS Management Console sign-in page. Sign in by using the account ID of the new member account and the Support1 IAM credentials. Set up the IAM users as required.

Answer: D

Explanation:

The best solution is to turn on the Concurrency Scaling feature for the Amazon Redshift cluster. This feature allows the cluster to automatically add additional capacity to handle bursts of read queries without affecting the performance of write queries. The additional capacity is transparent to the users and is billed separately based on the usage. This solution meets the business requirements of servicing read and write queries at all times and is also cost-effective compared to the other options, which involve provisioning additional resources or resizing the cluster. Reference: Amazon Redshift Documentation, Concurrency Scaling in Amazon Redshift

Question: 392

A data analytics company has an Amazon Redshift cluster that consists of several reserved nodes. The cluster 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 all 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 cluster. 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 cluster's CPU metrics in Amazon CloudWatch reach 80%.

- C. Deploy an AWS Lambda function to add capacity to the Amazon Redshift cluster by using an elastic resize operation when the cluster's CPU metrics in Amazon CloudWatch reach 80%.
- D. Turn on the Concurrency Scaling feature for the Amazon Redshift cluster.

Answer: C

Explanation:

The best solution is to deploy an AWS Lambda function to add capacity to the Amazon Redshift cluster by using an elastic resize operation when the cluster's CPU metrics in Amazon CloudWatch reach 80%. This solution will enable the cluster to scale up or down quickly by adding or removing nodes within minutes. This will improve the performance of the complex read queries and also reduce the cost by scaling down when the demand decreases. This solution is more cost-effective than using a classic resize operation, which takes longer and requires more downtime. It is also more suitable than using Amazon EMR, which is designed for big data processing rather than data warehousing. Reference: Amazon Redshift Documentation, Resizing clusters in Amazon Redshift, [Amazon EMR Documentation]

Question: 393

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:

The best solution is to deploy two firewall appliances into the shared services VPC, each in a separate

Availability Zone, and create a new Gateway Load Balancer to distribute traffic to them. A Gateway Load Balancer is designed for high performance and high availability scenarios with third-party network virtual appliances, such as firewalls. It operates at the network layer and maintains flow stickiness and symmetry to a specific appliance instance. It also uses the GENEVE protocol to encapsulate traffic between the load balancer and the appliances. To route traffic from other VPCs to the Gateway Load Balancer, a VPC Gateway Load Balancer endpoint is needed. This is a VPC endpoint that provides private connectivity between the appliances in the shared services VPC and the application servers in other VPCs. The endpoint must be added as the next hop in the route table for the shared services VPC. This solution ensures reliability and minimizes failover time between firewall appliances within a single AWS Region. Reference: What is a Gateway Load Balancer?, Gateway load balancer - Azure Load Balancer, Introducing Azure Gateway Load Balancer: Deploy and scale network ...

Question: 394

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:

The best solution is to use AWS Database Migration Service (AWS DMS) to migrate the data to AWS. AWS DMS is a web service that can migrate data from various sources to various targets, including MySQL databases. AWS DMS can perform full load and change data capture (CDC) migrations, which means that it can copy the existing data and also capture the ongoing changes to keep the source and target databases in sync. This minimizes the

downtime during the migration process. AWS DMS also supports encryption at rest and in transit by using AWS Key Management Service (AWS KMS) and TLS, respectively. This ensures that the data is protected during the migration. AWS DMS can also leverage AWS Direct Connect to transfer the data over a private connection, avoiding the internet. This solution meets all the requirements of the company. Reference: AWS Database Migration Service Documentation, Migrating Data to Amazon RDS for MySQL or MariaDB, Using SSL to Encrypt a Connection to a DB Instance

Question: 395 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 io2 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:

The best solution is to provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode and mount the EFS file system on each EC2 instance in the cluster. Amazon EFS is a fully managed, scalable, and elastic file storage service that supports the POSIX standard and can be accessed by multiple EC2 instances concurrently. Amazon EFS offers two performance modes: General Purpose and Max I/O. Max I/O mode is designed for highly parallelized workloads that can tolerate higher latencies than the General Purpose mode. Max I/O mode provides higher levels of aggregate throughput and operations per second, which are suitable for big data analytics applications. This solution meets all the requirements of the company. Reference: Amazon EFS Documentation, Amazon EFS performance modes

Question: 396

An online retail company is migrating its legacy on-premises .NET application to AWS. The application runs on load-balanced frontend web servers, load-balanced application servers, and a Microsoft SQL Server database.

The company wants to use AWS managed services where possible and does not want to rewrite the application. A solutions architect needs to implement a solution to resolve scaling issues and minimize licensing costs as the application scales.

Which solution will meet these requirements MOST cost-effectively?

- A. Deploy Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer for the web tier and for the application tier. Use Amazon Aurora PostgreSQL with BabelFish turned on to replatform the SQL Server database.
- B. Create images of all the servers by using AWS Database Migration Service (AWS DMS). Deploy Amazon EC2 instances that are based on the on-premises imports. Deploy the instances in an Auto Scaling group behind a Network Load Balancer for the web tier and for the application tier. Use Amazon DynamoDB as the database tier.
- C. Containerize the web frontend tier and the application tier. Provision an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Create an Auto Scaling group behind a Network Load Balancer for the web tier and for the application tier. Use Amazon RDS for SQL Server to host the database.
- D. Separate the application functions into AWS Lambda functions. Use Amazon API Gateway for the web frontend tier and the application tier. Migrate the data to Amazon S3. Use Amazon Athena to query the data.

Answer: A

Explanation:

The best solution is to create a tag policy that contains the allowed project tag values in the organization's management account and create an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added. A tag policy is a type of policy that can help standardize tags across resources in the organization's accounts. A tag policy can specify the allowed tag keys, values, and case treatment for compliance. A service control policy (SCP) is a type of policy that can restrict the actions that users and roles can perform in the organization's accounts. An SCP can deny access to specific API operations unless certain conditions are met, such as having a specific tag. By creating a tag policy in the management account and attaching it to each OU, the organization can enforce consistent tagging across all accounts. By creating an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added, the organization can prevent users from creating new resources without proper tagging. This solution will meet the requirements with the least effort, as it does not involve creating additional resources or modifying existing ones. Reference: Tag policies - AWS Organizations, Service control policies - AWS Organizations, AWS CloudFormation User Guide

Question: 397

A company uses an organization in AWS Organizations to manage the company's AWS accounts. The company uses AWS CloudFormation to deploy all infrastructure. A finance team wants to build a chargeback model. The finance team asked each business unit to tag resources by using a predefined list of project values. When the finance team used the AWS Cost and Usage Report in AWS Cost Explorer and filtered based on project, the team noticed noncompliant project values. The company wants to enforce the use of project tags for new resources.

Which solution will meet these requirements with the LEAST effort?

- A. Create a tag policy that contains the allowed project tag values in the organization's management account. Create an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added. Attach the SCP to each OU.
- B. Create a tag policy that contains the allowed project tag values in each OU. Create an SCP that denies the

cloudformation:CreateStack API operation unless a project tag is added. Attach the SCP to each OU.

C. Create a tag policy that contains the allowed project tag values in the AWS management account. Create an IAM policy that denies the cloudformation:CreateStack API operation unless a project tag is added. Assign the policy to each user.

D. Use AWS Service Catalog to manage the CloudFormation stacks as products. Use a TagOptions library to control project tag values. Share the portfolio with all OUs that are in the organization.

Answer: A

Explanation:

The best solution is to create a tag policy that contains the allowed project tag values in the organization's management account and create an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added. A tag policy is a type of policy that can help standardize tags across resources in the organization's accounts. A tag policy can specify the allowed tag keys, values, and case treatment for compliance. A service control policy (SCP) is a type of policy that can restrict the actions that users and roles can perform in the organization's accounts. An SCP can deny access to specific API operations unless certain conditions are met, such as having a specific tag. By creating a tag policy in the management account and attaching it to each OU, the organization can enforce consistent tagging across all accounts. By creating an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added, the organization can prevent users from creating new resources without proper tagging. This solution will meet the requirements with the least effort, as it does not involve creating additional resources or modifying existing ones. Reference: Tag policies - AWS Organizations, Service control policies - AWS Organizations, AWS CloudFormation User Guide

Question: 398

An online survey company runs its application in the AWS Cloud. The application is distributed and consists of microservices that run in an automatically scaled Amazon Elastic Container Service (Amazon ECS) cluster. The ECS cluster is a target for an Application Load Balancer (ALB). The ALB is a custom origin for an Amazon CloudFront distribution.

The company has a survey that contains sensitive data.

a. The sensitive data must be encrypted when it moves through the application. The application's data-handling microservice is the only microservice that should be able to decrypt the data. Which solution will meet these requirements?

A. Create a symmetric AWS Key Management Service (AWS KMS) key that is dedicated to the data-handling microservice. Create a field-level encryption profile and a configuration. Associate the KMS key and the configuration with the CloudFront cache behavior.

B. Create an RSA key pair that is dedicated to the data-handling microservice. Upload the public key to the CloudFront distribution. Create a field-level encryption profile and a configuration. Add the configuration to the CloudFront cache behavior.

C. Create a symmetric AWS Key Management Service (AWS KMS) key that is dedicated to the data-handling microservice. Create a Lambda@Edge function. Program the function to use the KMS key to encrypt the sensitive data.

D. Create an RSA key pair that is dedicated to the data-handling microservice. Create a

Lambda@Edge function. Program the function to use the private key of the RSA key pair to encrypt the sensitive data.

Answer: B

Explanation:

The best solution is to create an RSA key pair that is dedicated to the data-handling microservice and upload the public key to the CloudFront distribution. Then, create a field-level encryption profile and a configuration, and add the configuration to the CloudFront cache behavior. This solution will ensure that the sensitive data is encrypted at the edge locations of CloudFront, close to the end users, and remains encrypted throughout the application stack. Only the data-handling microservice, which has access to the private key of the RSA key pair, can decrypt the data. This solution does not require any additional resources or code changes, and leverages the built-in feature of CloudFront field-level encryption. For more information about CloudFront field-level encryption, see [Using fieldlevelencryption to help protect sensitive data](#).

Question: 399

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 THREE.)

- 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:

The best solution is to 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, and invoke an AWS Lambda function to perform image processing when a message is available in the queue. This solution will ensure that image processing can scale to handle the load, as Amazon S3 can store any amount of data and handle concurrent uploads, Amazon SQS can buffer the messages and deliver them reliably, and AWS Lambda can run code without provisioning or managing servers and scale automatically based on the demand. This solution will also notify the user when processing is complete by sending a push notification to the mobile app

using Amazon Simple Notification Service (Amazon SNS), which is a web service that enables applications to send and receive notifications from the cloud. This solution is more cost-effective than using Amazon MQ, which is a managed message broker service for Apache ActiveMQ that requires a dedicated broker instance, or S3 Batch Operations, which is a feature that allows users to perform bulk actions on S3 objects, such as copying or tagging, but does not support custom code execution. This solution is also more suitable than using Amazon Simple Email Service (Amazon SES), which is a web service that enables applications to send and receive email messages, but does not support push notifications for mobile devices. Reference: Amazon S3 Documentation, Amazon SQS Documentation, AWS Lambda Documentation, Amazon SNS Documentation

Question: 400

A company has mounted sensors to collect information about environmental parameters such as humidity and light throughout all the company's factories. The company needs to stream and analyze the data in the AWS Cloud in real time. If any of the parameters fall out of acceptable ranges, the factory operations team must receive a notification immediately.

Which solution will meet these requirements?

- A. Stream the data to an Amazon Kinesis Data Firehose delivery stream. Use AWS Step Functions to consume and analyze the data in the Kinesis Data Firehose delivery stream. use Amazon Simple Notification Service (Amazon SNS) to notify the operations team.
- B. Stream the data to an Amazon Managed Streaming for Apache Kafka (Amazon MSK) cluster. Set up a trigger in Amazon MSK to invoke an AWS Fargate task to analyze the data. Use Amazon Simple Email Service (Amazon SES) to notify the operations team.
- C. Stream the data to an Amazon Kinesis data stream. Create an AWS Lambda function to consume the Kinesis data stream and to analyze the data. Use Amazon Simple Notification Service (Amazon SNS) to notify the operations team.
- D. Stream the data to an Amazon Kinesis Data Analytics application. Use an automatically scaled and containerized service in Amazon Elastic Container Service (Amazon ECS) to consume and analyze the data. use Amazon Simple Email Service (Amazon SES) to notify the operations team.

Answer: C

Explanation:

The best solution is to stream the data to an Amazon Kinesis data stream and create an AWS Lambda function to consume the Kinesis data stream and to analyze the data. Amazon Kinesis is a web service that can collect, process, and analyze real-time streaming data from various sources, such as sensors. AWS Lambda is a serverless computing service that can run code in response to events, such as incoming data from a Kinesis data stream. By using AWS Lambda, the company can avoid provisioning or managing servers and scale automatically based on the demand. Amazon Simple Notification Service (Amazon SNS) is a web service that enables applications to send and receive notifications from the cloud. By using Amazon SNS, the company can notify the operations team immediately if any of the parameters fall out of acceptable ranges. This solution meets all the requirements of the company.

Reference: Amazon Kinesis Documentation, AWS Lambda Documentation, Amazon Simple Notification Service Documentation

Question: 401

A company is deploying AWS Lambda functions that access an Amazon RDS for PostgreSQL database. The company needs to launch the Lambda functions in a QA environment and in a production environment.

The company must not expose credentials within application code and must rotate passwords automatically.

Which solution will meet these requirements?

- A. Store the database credentials for both environments in AWS Systems Manager Parameter Store. Encrypt the credentials by using an AWS Key Management Service (AWS KMS) key. Within the application code of the Lambda functions, pull the credentials from the Parameter Store parameter by using the AWS SDK for Python (Boto3). Add a role to the Lambda functions to provide access to the Parameter Store parameter.
- B. Store the database credentials for both environments in AWS Secrets Manager with distinct key entry for the QA environment and the production environment. Turn on rotation. Provide a reference to the Secrets Manager key as an environment variable for the Lambda functions.
- C. Store the database credentials for both environments in AWS Key Management Service (AWS KMS). Turn on rotation. Provide a reference to the credentials that are stored in AWS KMS as an environment variable for the Lambda functions.
- D. Create separate S3 buckets for the QA environment and the production environment. Turn on server-side encryption with AWS KMS keys (SSE-KMS) for the S3 buckets. Use an object naming pattern that gives each Lambda function's application code the ability to pull the correct credentials for the function's corresponding environment. Grant each Lambda function's execution role access to Amazon S3.

Answer: B

Explanation:

The best solution is to store the database credentials for both environments in AWS Secrets Manager with distinct key entry for the QA environment and the production environment. AWS Secrets Manager is a web service that can securely store, manage, and retrieve secrets, such as database credentials. AWS Secrets Manager also supports automatic rotation of secrets by using Lambda functions or built-in rotation templates. By storing the database credentials for both environments in AWS Secrets Manager, the company can avoid exposing credentials within application code and rotate passwords automatically. By providing a reference to the Secrets Manager key as an environment variable for the Lambda functions, the company can easily access the credentials from the code by using the AWS SDK. This solution meets all the requirements of the company. Reference: AWS Secrets Manager Documentation, Using AWS Lambda with AWS Secrets Manager, Using environment variables - AWS Lambda

Question: 402

An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a solutions architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible for receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.

Which of the following is the MOST reliable approach to meet the requirements?

- A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.
- B. Receive the orders in an Amazon SQS queue and invoke an AWS Lambda function to process them.
- C. Receive the orders using the AWS Step Functions program and launch an Amazon ECS container to process them.
- D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

Answer: B

Explanation:

The best option is to use Amazon SQS and AWS Lambda to create a serverless order processing application. Amazon SQS is a fully managed message queue service that can decouple the order receiving and processing components, making the application more scalable and fault-tolerant. AWS Lambda is a serverless compute service that can automatically scale to handle the incoming messages from the SQS queue and process them according to the business logic. AWS Lambda can also integrate with Amazon DynamoDB to store the processed orders in a fast and flexible NoSQL database. This approach eliminates the need to provision, manage, or scale any servers or containers, and reduces the operational overhead and cost.

Option A is not reliable because using an EC2-hosted database to receive the orders introduces a single point of failure and a scalability bottleneck. EC2 instances also require more management and configuration than serverless services.

Option C is not reliable because using AWS Step Functions to receive the orders adds unnecessary complexity and cost to the application. AWS Step Functions is a service that coordinates multiple AWS services into a serverless workflow, but it is not designed to handle high-volume, sporadic, or unpredictable traffic patterns. AWS Step Functions also charges per state transition, which can be expensive for a large number of orders. Launching an ECS container to process each order also requires more resources and management than invoking a Lambda function.

Option D is not reliable because using Amazon Kinesis Data Streams to receive the orders is not suitable for this use case. Amazon Kinesis Data Streams is a service that enables real-time processing of streaming data at scale, but it is not meant for asynchronous message queuing. Amazon Kinesis Data Streams requires consumers to poll the data from the stream, which can introduce latency and complexity. Amazon Kinesis Data Streams also charges per shard hour, which can be expensive for a sporadic traffic pattern.

Amazon SQS

AWS Lambda

Amazon DynamoDB

AWS Step Functions Amazon ECS

Question: 403

A company maintains information on premises in approximately 1 million .csv files that are hosted on a VM. The data initially is 10 TB in size and grows at a rate of 1 TB each week. The company needs to **automate** backups of the data to the AWS Cloud.

Backups of the data must occur daily. The company needs a solution that applies custom filters to back up only a subset of the data that is located in designated source directories. The company has **set up an AWS Direct Connect** connection.

Which solution will meet the backup requirements with the **LEAST** operational overhead?

- A. Use the Amazon S3 CopyObject API operation with multipart upload to copy the existing data to Amazon S3. Use the CopyObject API operation to replicate new data to Amazon S3 daily.
- B. Create a backup plan in AWS Backup to back up the data to Amazon S3. Schedule the backup plan to **run daily**.
- C. Install the AWS DataSync agent as a VM that runs on the on-premises hypervisor. Configure a DataSync task to replicate the data to Amazon S3 daily.
- D. Use an AWS Snowball Edge device for the initial backup. Use AWS DataSync for incremental backups to Amazon S3 daily.

Answer: C

Explanation:

AWS DataSync is an online data transfer service that is designed to help customers get their data to and from AWS quickly, easily, and securely. Using DataSync, you can copy data from your on-premises NFS or SMB shares directly to Amazon S3, Amazon EFS, or Amazon FSx for Windows File Server. DataSync uses a purpose-built, parallel transfer protocol for speeds up to 10x faster than open source tools. DataSync also has built-in verification of data both in flight and at rest, so you can be confident that your data was transferred successfully. DataSync allows you to apply filters to select which files or folders to transfer, based on file name, size, or modification time. You can also schedule your DataSync tasks to run daily, weekly, or monthly, or on demand. DataSync is integrated with AWS Direct Connect, so you can take advantage of your existing private connection to AWS. DataSync is also a fully managed service, so you do not need to provision, configure, or maintain any infrastructure for data transfer.

Option A is incorrect because the Amazon S3 CopyObject API operation does not support filtering or scheduling, and it would require you to write and maintain custom scripts to automate the backup process.

Option B is incorrect because AWS Backup does not support filtering or transferring data from on-premises sources to Amazon S3. AWS Backup is a fully managed backup service that makes it easy to **centralize and automate** the backup of data across AWS services.

Option D is incorrect because AWS Snowball Edge is a physical device that is used for offline data transfer when network bandwidth is limited or unavailable. It is not suitable for daily backups or incremental transfers. AWS Snowball Edge also does not support filtering or scheduling.

1: Considering four different replication options for data in Amazon S3

2: Protect your file and backup archives using AWS DataSync and Amazon S3 Glacier

3: AWS DataSync FAQs

Question: 404

A company hosts an intranet web application on Amazon EC2 instances behind an Application Load Balancer (ALB). Currently, users authenticate to the application against an internal user database. The company needs to authenticate users to the application by using an existing AWS Directory Service for Microsoft Active Directory directory. All users with accounts in the directory must have access to the application. Which solution will meet these requirements?

- A. Create a new app client in the directory. Create a listener rule for the ALB. Specify the authenticate-oidc action for the listener rule. Configure the listener rule with the appropriate issuer, client ID and secret, and endpoint details for the Active Directory service. Configure the new app client with the callback URL that the ALB provides.
- B. Configure an Amazon Cognito user pool. Configure the user pool with a federated identity provider (IdP) that has metadata from the directory. Create an app client. Associate the app client with the user pool. Create a listener rule for the ALB. Specify the authenticate-cognito action for the listener rule. Configure the listener rule to use the user pool and app client.
- C. Add the directory as a new IAM identity provider (IdP). Create a new IAM role that has an entity type of SAML 2.0 federation. Configure a role policy that allows access to the ALB. Configure the new role as the default authenticated user role for the IdP. Create a listener rule for the ALB. Specify the authenticate-oidc action for the listener rule.
- D. Enable AWS IAM Identity Center (AWS Single Sign-On). Configure the directory as an external identity provider (IdP) that uses SAML. Use the automatic provisioning method. Create a new IAM role that has an entity type of SAML 2.0 federation. Configure a role policy that allows access to the ALB. Attach the new role to all groups. Create a listener rule for the ALB. Specify the authenticate-cognito action for the listener rule.

Answer: A

Explanation:

The correct solution is to use the authenticate-oidc action for the ALB listener rule and configure it with the details of the AWS Directory Service for Microsoft Active Directory directory. This way, the ALB can use OpenID Connect (OIDC) to authenticate users against the directory and grant them access to the intranet web application. The app client in the directory is used to register the ALB as an OIDC client and provide the necessary credentials and endpoints. The callback URL is the URL that the ALB redirects the user to after a successful authentication. This solution does not require any additional services or roles, and it leverages the existing directory accounts for all users.

The other solutions are incorrect because they either use the wrong action for the ALB listener rule, or they involve unnecessary or incompatible services or roles. For example:

Solution B is incorrect because it uses Amazon Cognito user pool, which is a separate user directory service that does not integrate with AWS Directory Service for Microsoft Active Directory. To use this solution, the company would have to migrate or synchronize their users from the directory to the user pool, which is not required by the question. Moreover, the authenticate-cognito action for the ALB listener rule only works with Amazon Cognito user pools, not with federated identity providers (IdPs) that have metadata from the directory. Solution C is incorrect because it uses IAM as an identity provider (IdP), which is not compatible with AWS Directory Service for Microsoft Active Directory. IAM can only be used as an IdP for web identity federation, which allows users to sign in with social media or other third-party IdPs, not with Active Directory. Moreover,

the authenticate-oidc action for the ALB listener rule requires an OIDC IdP, not a SAML 2.0 federation IdP, which is what IAM provides.

Solution D is incorrect because it uses AWS IAM Identity Center (AWS Single Sign-On), which is a service that simplifies the management of SSO access to multiple AWS accounts and business applications. This service is not needed for the scenario in the question, which only involves a single intranet web application. Moreover, the authenticate-cognito action for the ALB listener rule does not work with external IdPs that use SAML, such as AWS IAM Identity Center.

Authenticate users using an Application Load Balancer

What is AWS Directory Service for Microsoft Active Directory?

Using OpenID Connect for user authentication

Question: 405

A company is launching a new online game on Amazon EC2 instances. The game must be available globally. The company plans to run the game in three AWS Regions: us-east-1, eu-west-1, and ap-southeast-1. The game's leaderboards, player inventory, and event status must be available across Regions.

A solutions architect must design a solution that will give any Region the ability to scale to handle the load of all Regions. Additionally, users must automatically connect to the Region that provides the least latency.

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

A. Create an EC2 Spot Fleet. Attach the Spot Fleet to a Network Load Balancer (NLB) in each Region. Create an AWS Global Accelerator IP address that points to the NLB. Create an Amazon Route 53 latency-based routing entry for the Global Accelerator IP address. Save the game metadata to an Amazon RDS for MySQL DB instance in each Region. Set up a read replica in the other Regions.

B. Create an Auto Scaling group for the EC2 instances. Attach the Auto Scaling group to a Network Load Balancer (NLB) in each Region. For each Region, create an Amazon Route 53 entry that uses geoproximity routing and points to the NLB in that Region. Save the game metadata to MySQL databases on EC2 instances in each Region. Save the game metadata to MySQL databases on EC2 instances in each Region. Set up replication between the database EC2 instances in each Region.

C. Create an Auto Scaling group for the EC2 instances. Attach the Auto Scaling group to a Network Load Balancer (NLB) in each Region. For each Region, create an Amazon Route 53 entry that uses latency-based routing and points to the NLB in that Region. Save the game metadata to an Amazon DynamoDB global table.

D. Use EC2 Global View. Deploy the EC2 instances to each Region. Attach the instances to a Network Load Balancer (NLB). Deploy a DNS server on an EC2 instance in each Region. Set up custom logic on each DNS server to redirect the user to the Region that provides the lowest latency. Save the game metadata to an Amazon Aurora global database.

Answer: C

Explanation:

The best option is to use an Auto Scaling group, a Network Load Balancer, Amazon Route 53, and Amazon DynamoDB to create a scalable, highly available, and low-latency online game application. An Auto Scaling group can automatically adjust the number of EC2 instances based on the demand and traffic in each Region. A

Network Load Balancer can distribute the incoming traffic across the EC2 instances and handle millions of requests per second. Amazon Route 53 can use latency-based routing to direct the users to the Region that provides the best performance. Amazon DynamoDB global tables can replicate the game metadata across multiple Regions, ensuring consistency and availability of the data. This approach minimizes the operational overhead and cost, as it leverages fully managed services and avoids custom logic or replication. Option A is not optimal because using an EC2 Spot Fleet can introduce the risk of losing the EC2 instances if the Spot price exceeds the bid price, which can affect the availability and performance of the game. Using AWS Global Accelerator can improve the network performance, but it is not necessary if Amazon Route 53 can already route the users to the closest Region. Using Amazon RDS for MySQL can store the game metadata, but it requires setting up read replicas and managing the replication lag across Regions, which can increase the complexity and cost.

Option B is not optimal because using geoproximity routing can direct the users to the Region that is geographically closer, but it does not account for the network latency or performance. Using MySQL databases on EC2 instances can store the game metadata, but it requires managing the EC2 instances, the database software, the backups, the patches, and the replication across Regions, which can increase the operational overhead and cost.

Option D is not optimal because using EC2 Global View is not a valid service. Using a custom DNS server on an EC2 instance can redirect the user to the Region that provides the lowest latency, but it requires developing and maintaining the custom logic, as well as managing the EC2 instance, which can increase the operational overhead and cost. Using Amazon Aurora global database can store the game metadata, but it is more expensive and complex than using Amazon DynamoDB global tables.

Auto Scaling groups Network Load Balancer Amazon Route 53 Amazon DynamoDB global tables

Question: 406

A solutions architect is preparing to deploy a new security tool into several previously unused AWS Regions. The solutions architect will deploy the tool by using an AWS CloudFormation stack set. The stack set's template contains an IAM role that has a custom name. Upon creation of the stack set, no stack instances are created successfully.

What should the solutions architect do to deploy the stacks successfully?

- A. Enable the new Regions in all relevant accounts. Specify the `CAPABILITY_NAMED_IAM` capability during the creation of the stack set.
- B. Use the Service Quotas console to request a quota increase for the number of CloudFormation stacks in each new Region in all relevant accounts. Specify the `CAPABILITY_IAM` capability during the creation of the stack set.
- C. Specify the `CAPABILITY_NAMED_IAM` capability and the `SELF_MANAGED` permissions model during the creation of the stack set.
- D. Specify an administration role ARN and the `CAPABILITY_IAM` capability during the creation of the stack set.

Answer: A

Explanation:

The CAPABILITY_NAMED_IAM capability is required when creating or updating CloudFormation stacks that contain IAM resources with custom names. This capability acknowledges that the template might create IAM resources that have broad permissions or affect other resources in the AWS account. The stack set's template contains an IAM role that has a custom name, so this capability is needed. Enabling the new Regions in all relevant accounts is also necessary to deploy the stack set across multiple Regions and accounts. Option B is incorrect because the Service Quotas console is used to view and manage the quotas for AWS services, not for CloudFormation stacks. The number of stacks per Region per account is not a service quota that can be increased.

Option C is incorrect because the SELF_MANAGED permissions model is used when the administrator wants to retain full permissions to manage stack sets and stack instances. This model does not affect the creation of the stack set or the requirement for the CAPABILITY_NAMED_IAM capability.

Option D is incorrect because an administration role ARN is optional when creating a stack set. It is used to specify a role that CloudFormation assumes to create stack instances in the target accounts. It does not affect the creation of the stack set or the requirement for the CAPABILITY_NAMED_IAM capability.

- 1: AWS CloudFormation stack sets
- 2: Acknowledging IAM resources in AWS CloudFormation templates
- 3: AWS CloudFormation stack set permissions

Question: 407

A company wants to design a disaster recovery (DR) solution for an application that runs in the company's data center. The application writes to an SMB file share and creates a copy on a second file share. Both file shares are in the data center. The application uses two types of files: metadata files and image files.

The company wants to store the copy on AWS. The company needs the ability to use SMB to access the data from either the data center or AWS if a disaster occurs. The copy of the data is rarely accessed but must be available within 5 minutes.

Which solution will meet these requirements MOST cost-effectively?

- A. Deploy AWS Outposts with Amazon S3 storage. Configure a Windows Amazon EC2 instance on Outposts as a file server.
- B. Deploy an Amazon FSx File Gateway. Configure an Amazon FSx for Windows File Server Multi-AZ file system that uses SSD storage.
- C. Deploy an Amazon S3 File Gateway. Configure the S3 File Gateway to use Amazon S3 Standard- Infrequent Access (S3 Standard-IA) for the metadata files and to use S3 Glacier Deep Archive for the image files.
- D. Deploy an Amazon S3 File Gateway. Configure the S3 File Gateway to use Amazon S3 Standard- Infrequent Access (S3 Standard-IA) for the metadata files and image files.

Answer: C

Explanation:

The correct solution is to use an Amazon S3 File Gateway to store the copy of the SMB file share on AWS. An S3 File Gateway enables on-premises applications to store and access objects in Amazon S3 using the SMB protocol. The S3 File Gateway can also be accessed from AWS using the SMB protocol, which provides the ability to use the data from either the data center or AWS if a disaster occurs. The S3 File Gateway

supports tiering of data to different S3 storage classes based on the file type. This allows the company to optimize the storage costs by using S3 Standard-Infrequent Access (S3 Standard-IA) for the metadata files, which are rarely accessed but must be available within 5 minutes, and S3 Glacier Deep Archive for the image files, which are the lowest-cost storage class and suitable for long-term retention of data that is rarely accessed. This solution is the most cost-effective because it does not require any additional hardware, software, or replication services.

The other solutions are incorrect because they either use more expensive or unnecessary services or components, or they do not meet the requirements. For example:

Solution A is incorrect because it uses AWS Outposts with Amazon S3 storage, which is a very expensive and complex solution for the scenario in the question. AWS Outposts is a service that extends AWS infrastructure, services, APIs, and tools to virtually any data center, co-location space, or on-premises facility. It is designed for customers who need low latency and local data processing. Amazon S3 storage on Outposts provides a subset of S3 features and APIs to store and retrieve data on Outposts. However, this solution does not provide SMB access to the data on Outposts, which requires a Windows EC2 instance on Outposts as a file server. This adds more cost and complexity to the solution, and it does not provide the ability to access the data from AWS if a disaster occurs. Solution B is incorrect because it uses Amazon FSx File Gateway and Amazon FSx for Windows File Server Multi-AZ file system that uses SSD storage, which are both more expensive and unnecessary services for the scenario in the question. Amazon FSx File Gateway is a service that enables on-premises applications to store and access data in Amazon FSx for Windows File Server using the SMB protocol. Amazon FSx for Windows File Server is a fully managed service that provides native Windows file shares with the compatibility, features, and performance that Windows-based applications rely on. However, this solution does not meet the requirements because it does not provide the ability to use different storage classes for the metadata files and image files, and it does not provide the ability to access the data from AWS if a disaster occurs. Moreover, using a Multi-AZ file system that uses SSD storage is overprovisioned and costly for the scenario in the question, which involves rarely accessed data that must be available within 5 minutes.

Solution D is incorrect because it uses an S3 File Gateway that uses S3 Standard-IA for both the metadata files and image files, which is not the most cost-effective solution for the scenario in the question. S3 Standard-IA is a storage class that offers high durability, availability, and performance for infrequently accessed data. However, it is more expensive than S3 Glacier Deep Archive, which is the lowest-cost storage class and suitable for long-term retention of data that is rarely accessed. Therefore, using S3 Standard-IA for the image files, which are likely to be larger and more numerous than the metadata files, is not optimal for the storage costs.

What is S3 File Gateway?

Using Amazon S3 storage classes with S3 File Gateway

Accessing your file shares from AWS

Question: 408

A company wants to migrate an Amazon Aurora MySQL DB cluster from an existing AWS account to a new AWS account in the same AWS Region. Both accounts are members of the same organization in AWS

Organizations.

The company must minimize database service interruption before the company performs DNS cutover to the new database.

Which migration strategy will meet this requirement?

- A. Take a snapshot of the existing Aurora database. Share the snapshot with the new AWS account. Create an Aurora DB cluster in the new account from the snapshot.
- B. Create an Aurora DB cluster in the new AWS account. Use AWS Database Migration Service (AWS DMS) to migrate data between the two Aurora DB clusters.
- C. Use AWS Backup to share an Aurora database backup from the existing AWS account to the new AWS account. Create an Aurora DB cluster in the new AWS account from the snapshot.
- D. Create an Aurora DB cluster in the new AWS account. Use AWS Application Migration Service to migrate data between the two Aurora DB clusters.

Answer: B

Explanation:

The best migration strategy to meet the requirement of minimizing database service interruption before the DNS cutover is to use AWS DMS to migrate data between the two Aurora DB clusters. AWS DMS can perform continuous replication of data with high availability and consolidate databases into a petabyte-scale data warehouse by streaming data to Amazon Redshift and Amazon S3¹. AWS DMS supports homogeneous migrations such as migrating from one Aurora MySQL DB cluster to another, as well as heterogeneous migrations between different database platforms². AWS DMS also supports cross-account migrations, as long as the source and target databases are in the same AWS Region³.

The other options are not optimal for the following reasons:

Option A: Taking a snapshot of the existing Aurora database and restoring it in the new account would require a downtime during the snapshot and restore process, which could be significant for large databases. Moreover, any changes made to the source database after the snapshot would not be replicated to the target database, resulting in data inconsistency⁴.

Option C: Using AWS Backup to share an Aurora database backup from the existing AWS account to the new AWS account would have the same drawbacks as option A, as AWS Backup uses snapshots to create backups of Aurora databases.

Option D: Using AWS Application Migration Service to migrate data between the two Aurora DB clusters is not a valid option, as AWS Application Migration Service is designed to migrate applications, not databases, to AWS. AWS Application Migration Service can migrate applications from on-premises or other cloud environments to AWS, using agentless or agent-based methods. **Reference:**

1: What Is AWS Database Migration Service? - AWS Database Migration Service

2: Sources for Data Migration - AWS Database Migration Service

3: AWS Database Migration Service FAQs

4: Working with DB Cluster Snapshots - Amazon Aurora

: [Backing Up and Restoring an Amazon Aurora DB Cluster - Amazon Aurora]

: [What is AWS Application Migration Service? - AWS Application Migration Service]

Question: 409

A company is planning to migrate an on-premises data center to AWS. The company currently hosts the data center on Linux-based VMware VMs. A solutions architect must collect information about network dependencies between the VMs. The information must be in the form of a diagram that details host IP addresses, hostnames, and network connection information.

Which solution will meet these requirements?

- A. Use AWS Application Discovery Service. Select an AWS Migration Hub home AWS Region. Install the AWS Application Discovery Agent on the on-premises servers for data collection. Grant permissions to Application Discovery Service to use the Migration Hub network diagrams.
- B. Use the AWS Application Discovery Service Agentless Collector for server data collection. Export the network diagrams from the AWS Migration Hub in .png format.
- C. Install the AWS Application Migration Service agent on the on-premises servers for data collection. Use AWS Migration Hub data in Workload Discovery on AWS to generate network diagrams.
- D. Install the AWS Application Migration Service agent on the on-premises servers for data collection. Export data from AWS Migration Hub in .csv format into an Amazon CloudWatch dashboard to generate network diagrams.

Answer: A

Explanation:

To effectively gather information about network dependencies between VMs in an on-premises data center for migration to AWS, it's crucial to use tools that can capture detailed application and server dependencies. The AWS Application Discovery Service is designed for this purpose, particularly when migrating from environments like Linux-based VMware VMs. By installing the AWS Application Discovery Agent on the on-premises servers, the service can collect necessary data such as host IP addresses, hostnames, and network connection information. This data is crucial for creating a comprehensive network diagram that outlines the interactions and dependencies between various components of the on-premises infrastructure. The integration with AWS Migration Hub enhances this process by allowing the visualization of these dependencies in a network diagram format, aiding in the planning and execution of the migration process. This approach ensures a thorough understanding of the on-premises environment, which is essential for a successful migration to AWS.

Reference:

AWS Documentation on Application Discovery Service: This provides detailed guidance on how to use the Application Discovery Service, including the installation and configuration of the Discovery Agent.

AWS Migration Hub User Guide: Offers insights on how to integrate Application Discovery Service data with Migration Hub for comprehensive migration planning and tracking.

AWS Solutions Architect Professional Learning Path: Contains advanced topics and best practices for migrating complex on-premises environments to AWS, emphasizing the use of AWS services and tools for effective migration planning and execution.

Question: 410

A company is planning a migration from an on-premises data center to the AWS cloud. The company plans to use multiple AWS accounts that are managed in an organization in AWS organizations. The company will cost a small number of accounts initially and will add accounts as needed. A solution architect must design a solution that turns on AWS accounts.

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

- A. Create an AWS Lambda function that creates a new cloudTrail trail in all AWS account in the organization.

Invoke the Lambda function daily by using a scheduled action in Amazon EventBridge.

- B. Create a new CloudTrail trail in the organizations management account. Configure the trail to log all events for all AYS accounts in the organization.
- C. Create a new CloudTrail trail in all AWS accounts in the organization. Create new trails whenever a new account is created.
- D. Create an AWS systems Manager Automaton runbook that creates a cloud trail in all AWS accounts in the organization. Invoke the automation by using Systems Manager State Manager.

Answer: B

Explanation:

The most operationally efficient solution for turning on AWS CloudTrail across multiple AWS accounts managed within an AWS Organization is to create a single CloudTrail trail in the organization's management account and configure it to log events for all accounts within the organization. This approach leverages CloudTrail's ability to consolidate logs from all accounts in an organization, thereby simplifying management, reducing overhead, and ensuring consistent logging across accounts. This method eliminates the need for manual intervention in each account, making it an operationally efficient choice for organizations planning to scale their AWS usage.

AWS CloudTrail Documentation: Provides detailed instructions on setting up CloudTrail, including how to configure it for an organization.

AWS Organizations Documentation: Offers insights into best practices for managing multiple AWS accounts and how services like CloudTrail integrate with AWS Organizations.

AWS Best Practices for Security and Governance: Guides on how to effectively use AWS services to maintain a secure and well-governed AWS environment, with a focus on centralized logging and monitoring.

Question: 411

A company's factory and automaton applications are running in a single VPC. More than 23 applications run on a combination of Amazon EC2, Amazon Elastic Container Service (Amazon ECS), and Amazon RDS.

The company has software engineers spread across three teams. One of the three teams owns each application, and each team is responsible for the cost and performance of all of its applications. Team resources have tags that represent their application and team. The teams use IAM access for daily activities.

The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs item the last 12 months and to help forecast costs for the next 12 months. A solution architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.

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

- A. Activate the user-defined cost allocation tags that represent the application and the team.
- B. Activate the AWS generated cost allocation tags that represent the application and the team.
- C. Create a cost category for each application in Billing and Cost Management
- D. Activate IAM access to Billing and Cost Management.
- E. Create a cost budget
- F. Enable Cost Explorer.

Answer: A,C,F

Explanation:

To attribute AWS costs to specific applications or teams and enable detailed cost analysis and forecasting, the solution architect should recommend the following actions: A. Activating user-defined cost allocation tags for resources associated with each application and team allows for detailed tracking of costs by these identifiers. C. Creating a cost category for each application within AWS Billing and Cost Management enables the organization to group costs according to application, facilitating detailed reporting and analysis. F. Enabling Cost Explorer is essential for analyzing and visualizing AWS spending over time. It provides the capability to view historical costs and forecast future expenses, supporting the company's requirement for cost comparison and forecasting.

AWS Billing and Cost Management Documentation: Covers the activation of cost allocation tags, creation of cost categories, and the use of Cost Explorer for cost management.

AWS Tagging Strategies: Provides best practices for implementing tagging strategies that support cost allocation and reporting.

AWS Cost Explorer Documentation: Details how to use Cost Explorer to analyze and forecast AWS costs.

Question: 412

A company is using an organization in AWS organization to manage AWS accounts. For each new project the company creates a new linked account. After the creation of a new account, the root user signs in to the new account and creates a service request to increase the service quota for Amazon EC2 instances. A solutions architect needs to automate this process.

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

- A. Create an Amazon EventBridge rule to detect creation of a new account. Send the event to an Amazon Simple Notification Service (Amazon SNS) topic that invokes an AWS Lambda function. Configure the Lambda function to run the request-service-quota-increase command to request a service quota increase for EC2 instances.
- B. Create a Service Quotas request template in the management account. Configure the desired service quota increases for EC2 instances.
- C. Create an AWS Config rule in the management account to set the service quota for EC2 instances.
- D. Create an Amazon EventBridge rule to detect creation of a new account. Send the event to an Amazon Simple Notification Service (Amazon SNS) topic that involves an AWS Lambda function. Configure the Lambda function to run the create-case command to request a service quota increase for EC2 instances.

Answer: A

Explanation:

Automating the process of increasing service quotas for Amazon EC2 instances in new AWS accounts with minimal operational overhead can be effectively achieved by using Amazon EventBridge, Amazon SNS, and AWS Lambda. An EventBridge rule can detect the creation of a new account and trigger an SNS topic, which in turn invokes a Lambda function. This function can then programmatically request a service quota

increase for EC2 instances using the AWS Service Quotas API. This approach streamlines the process, reduces manual intervention, and ensures that new accounts are automatically configured with the desired service quotas.

Amazon EventBridge Documentation: Provides guidance on setting up event rules for detecting AWS account creation.

AWS Lambda Documentation: Details how to create and configure Lambda functions to perform automated tasks, such as requesting service quota increases.

AWS Service Quotas Documentation: Offers information on managing and requesting increases for AWS service quotas programmatically.

Question: 413 A company has Linux-based Amazon EC2 instances. Users must access the instances by using SSH with EC2 SSH Key pairs. Each machine requires a unique EC2 Key pair.

The company wants to implement a key rotation policy that will, upon request, automatically rotate all the EC2 key pairs and keep the key in a securely encrypted place. The company will accept less than 1 minute of downtime during key rotation.

Which solution will meet these requirements?

- A. Store all the keys in AWS Secrets Manager. Define a Secrets Manager rotation schedule to invoke an AWS Lambda function to generate new key pairs. Replace public keys on EC2 instances. Update the private keys in Secrets Manager.
- B. Store all the keys in Parameter Store. Store, a capability of AWS Systems Manager, as a string. Define a Systems Manager maintenance window to invoke an AWS Lambda function to generate new key pairs. Replace public keys on EC2 instance. Update the private keys in parameter.
- C. Import the EC2 key pairs into AWS Key Management Service (AWS KMS). Configure automatic key rotation for these key pairs. Create an Amazon EventBridge scheduled rule to invoke an AWS Lambda function to initiate the key rotation AWS KMS.
- D. Add all the EC2 instances to Fleet Manager, a capability of AWS Systems Manager. Define a Systems Manager maintenance window to issue a Systems Manager Run Command document to generate new Key pairs and to rotate public keys to all the instances in Fleet Manager.

Answer: A

Explanation:

To meet the requirements for automatic key rotation of EC2 SSH key pairs with minimal downtime, storing the keys in AWS Secrets Manager and defining a rotation schedule is the most suitable solution. AWS Secrets Manager supports automatic rotation of secrets, including SSH keys, by invoking a Lambda function that can handle the creation of new key pairs and the replacement of public keys on EC2 instances. Updating the corresponding private keys in Secrets Manager ensures secure and centralized management of SSH keys, complying with the key rotation policy and minimizing operational overhead.

AWS Secrets Manager Documentation: Describes how to store and rotate secrets, including SSH keys, using Secrets Manager and Lambda functions.

AWS Lambda Documentation: Provides information on creating Lambda functions for custom secret rotation logic.

AWS Best Practices for Security: Highlights the importance of key rotation and how AWS services like Secrets Manager can facilitate secure and automated key management.

Question: 414

A company has many services running in its on-premises data center. The data center is connected to AWS using AWS Direct Connect (DX) and an IPsec VPN. The service data is sensitive and connectivity cannot traverse the internet. The company wants to expand to a new market segment and begin offering its services to other companies that are using AWS.

Which solution will meet these requirements?

- A. Create a VPC Endpoint Service that accepts TCP traffic, host it behind a Network Load Balancer, and make the service available over DX.
- B. Create a VPC Endpoint Service that accepts HTTP or HTTPS traffic, host it behind an Application Load Balancer, and make the service available over DX.
- C. Attach an internet gateway to the VPC. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.
- D. Attach a NAT gateway to the VPC. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.

Answer: B

Explanation:

To offer services to other companies using AWS without traversing the internet, creating a VPC Endpoint Service hosted behind an Application Load Balancer (ALB) and making it available over AWS Direct Connect (DX) is the most suitable solution. This approach ensures that the service traffic remains within the AWS network, adhering to the requirement that connectivity must not traverse the internet. An ALB is capable of handling HTTP/HTTPS traffic, making it appropriate for web-based services. Utilizing DX for connectivity between the on-premises data center and AWS further secures and optimizes the network path.

AWS Direct Connect Documentation: Explains how to set up DX for private connectivity between AWS and an on-premises network.

Amazon VPC Endpoint Services (AWS PrivateLink) Documentation: Provides details on creating and configuring endpoint services for private, secure access to services hosted in AWS.

AWS Application Load Balancer Documentation: Offers guidance on configuring ALBs to distribute HTTP/HTTPS traffic efficiently.

Question: 415

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 with less than 500 milliseconds or latency to end users. The solution also must be highly available and must have an RTO of 30 seconds.

Which solution will meet these requirements?

- A. Use an Amazon Redshift cluster to store the data. Use a state website that is hosted on Amazon S3 with backend APIs that are 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:

For an application requiring low-latency access to financial information and high availability with a Recovery Time Objective (RTO) of 30 seconds, using Amazon DynamoDB for data storage and Amazon QuickSight for reporting is the most suitable solution. DynamoDB offers fast, consistent, and single-digit millisecond latency for data retrieval, meeting the latency requirements. QuickSight's ability to directly query DynamoDB datasets and provide embedded dashboards for reporting enables real-time financial report generation. This combination ensures high availability and meets the RTO requirement, providing a robust solution for the application's needs.

Amazon DynamoDB Documentation: Describes the features and benefits of DynamoDB, emphasizing its performance and scalability for applications requiring low-latency access to data.

Amazon QuickSight Documentation: Provides information on using QuickSight for creating and embedding interactive dashboards, including direct querying of DynamoDB datasets for real-time data visualization.

Question: 416

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 assets become unavailable in the closest Region, they should be fetched from the other Region.

What should a solutions architect do to meet these requirements?

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

Answer: A

Explanation:

To ensure that game assets are fetched from the closest region and have a fallback option in case the assets become unavailable in the closest region, a solution architect should leverage Amazon CloudFront, a global content delivery network (CDN) service. By creating an Amazon CloudFront distribution and setting up origin groups, the architect can specify multiple origins (in this case, the Application Load Balancers in each region). The primary origin will serve content under normal circumstances, and if the content becomes unavailable, CloudFront will automatically switch to the secondary origin. This approach not only meets the requirement of regional proximity and redundancy but also optimizes latency and enhances the gaming experience by serving assets from the nearest geographical location to the end-user.

AWS Documentation on Amazon CloudFront and origin groups provides detailed instructions on setting up distributions with multiple origins for high availability and performance optimization. Additionally, AWS whitepapers and best practices on content delivery and global applications offer insights into effectively utilizing CloudFront and other AWS services to achieve low latency and high availability.

Question: 417

A company use an organization in AWS Organizations to manage multiple AWS accounts. The company hosts some applications in a VPC in the company's snared services account. The company has attached a transit gateway to the VPC in the Shared services account.

The company is developing a new capability and has created a development environment that requires access to the applications that are in the snared services account. The company intends to delete and recreate resources frequently in the development account. The company also wants to give a development team the ability to recreate the team's connection to the shared services account as required.

Which solution will meet these requirements?

- A. Create a transit gateway in the development account. Create a transit gateway peering request to the shared services account. Configure the snared services transit gateway to automatically accept peering connections.
- B. Turn on automate acceptance for the transit gateway in the shared services account. Use AWS Resource Access Manager (AWS RAM) to share the transit gateway resource in the shared services account with the development account. Accept the resource in tie development account. Create a transit gateway attachment in the development account.
- C. Turn on automate acceptance for the transit gateway in the shared services account. Create a VPC endpoint. Use the endpoint policy to grant permissions on the VPC endpoint for the development account. Configure the endpoint service to automatically accept connection requests. Provide the endpoint details to the development team.
- D. Create an Amazon EventBridge rule to invoke an AWS Lambda function that accepts the transit gateway attachment value the development account makes an attachment request. Use AWS Network Manager to store. The transit gateway in the shared services account with the development account. Accept the transit gateway in the development account.

Answer: B

Explanation:

For a development environment that requires frequent resource recreation and connectivity to applications hosted in a shared services account, the most efficient solution involves using AWS Resource Access Manager (RAM) and the transit gateway in the shared services account. By turning on automatic acceptance for the transit gateway in the shared services account and sharing it with the development account through AWS RAM, the development team can easily recreate their connection as needed without manual intervention. This setup allows for scalable, flexible connectivity between accounts while minimizing operational overhead and ensuring consistent access to shared services.

AWS Documentation on AWS Resource Access Manager and Transit Gateway provides guidance on sharing network resources across AWS accounts and enabling automatic acceptance for transit gateway attachments. This approach is also supported by AWS best practices for multi-account strategies using AWS Organizations and network architecture.

Question: 418

A company uses AWS Organizations AWS account. A solutions architect must design a solution in which only administrator roles are allowed to use IAM actions. However the solutions architect does not have access to all the AWS account throughout the company.

Which solution meets these requirements with the LEAST operational overhead?

- A. Create an SCP that applies to all the AWS accounts to allow IAM actions only for administrator roles. Apply the SCP to the root OLI.
- B. Configure AWS CloudTrail to invoke an AWS Lambda function for each event that is related to IAM actions. Configure the function to deny the action. If the user who invoked the action is not an administrator.
- C. Create an SCP that applies to all the AWS accounts to deny IAM actions for all users except for those with administrator roles. Apply the SCP to the root OU.
- D. Set an IAM permissions boundary that allows IAM actions. Attach the permissions boundary to every administrator role across all the AWS accounts.

Answer: A

Explanation:

To restrict IAM actions to only administrator roles across all AWS accounts in an organization, the most operationally efficient solution is to create a Service Control Policy (SCP) that allows IAM actions exclusively for administrator roles and apply this SCP to the root Organizational Unit (OU) of AWS Organizations. This method ensures a centralized governance mechanism that uniformly applies the policy across all accounts, thereby minimizing the need for individual account-level configurations and reducing operational complexity. AWS Documentation on AWS Organizations and Service Control Policies offers comprehensive information on creating and managing SCPs for organizational-wide policy enforcement. This approach aligns with AWS best practices for managing permissions and ensuring secure and compliant account configurations within an AWS Organization.

Question: 419

A company needs to migrate an on-premises SFTP site to AWS. The SFTP site currently runs on a Linux VM. Uploaded files are made available to downstream applications through an NFS share. As part of the migration to AWS, a solutions architect must implement high availability. The solution must provide external vendors with a set of static public IP addresses that the vendors can allow. The company has set up an AWS Direct Connect connection between its on-premises data center and its VPC.

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

- A. Create an AWS Transfer Family server, configure an internet-facing VPC endpoint for the Transfer Family server, specify an Elastic IP address for each subnet, configure the Transfer Family server to place files into an Amazon Elastic File System (Amazon EFS) file system that is deployed across multiple Availability Zones. Modify the configuration on the downstream applications that access the existing NFS share to mount the EFS endpoint instead.
- B. Create an AWS Transfer Family server. Configure a publicly accessible endpoint for the Transfer Family server. Configure the Transfer Family server to place files into an Amazon Elastic File System [Amazon EFS] the system that is deployed across multiple Availability Zones. Modify the configuration on the downstream applications that access the existing NFS share to mount the its endpoint instead.
- C. Use AWS Application Migration service to migrate the existing Linux VM to an Amazon EC2 instance. Assign an Elastic IP address to the EC2 instance. Mount an Amazon Elastic File system (Amazon EFS) the system to the EC2 instance. Configure the SFTP server to place files in. the EFSfile system. Modify the configuration on the downstream applications that access the existing NFS share to mount the EFS endpoint instead.
- D. Use AWS Application Migration Service to migrate the existing Linux VM to an AWS Transfer Family server. Configure a publicly accessible endpoint for the Transfer Family server. Configure the Transfer Family sever to place files into an Amazon FSx for Luster the system that is deployed across multiple Availability Zones. Modify the configuration on the downstream applications that access the existing NFS share to mount the FSx for Luster endpoint instead.

Answer: A

Explanation:

To migrate an on-premises SFTP site to AWS with high availability and a set of static public IP addresses for external vendors, the best solution is to create an AWS Transfer Family server with an internet-facing VPC endpoint. Assigning Elastic IP addresses to each subnet and configuring the server to store files in an Amazon Elastic File System (EFS) that spans multiple Availability Zones ensures high availability and consistent access. This approach minimizes operational overhead by leveraging AWS managed services and eliminates the need to manage underlying infrastructure. AWS Documentation on AWS Transfer Family and Amazon Elastic File System provides detailed instructions on setting up a highly available SFTP environment on AWS. This solution is in line with AWS best practices for migrating and modernizing applications with minimal disruption and ensuring high availability and security.

Question: 420

A company's compliance audit reveals that some Amazon Elastic Block Store (Amazon EBS) volumes that were

created in an AWS account were not encrypted. A solutions architect must implement a solution to encrypt all new EBS volumes at rest

Which solution will meet this requirement with the LEAST effort?

- A. Create an Amazon EventBridge rule to detect the creation of unencrypted EBS volumes. Invoke an AWS Lambda function to delete noncompliant volumes.
- B. Use AWS Audit Manager with data encryption.
- C. Create an AWS Config rule to detect the creation of a new EBS volume. Encrypt the volume by using AWS Systems Manager Automation.
- D. Turn on EBS encryption by default in all AWS Regions.

Answer: D

Explanation:

The most effortless way to ensure that all new Amazon Elastic Block Store (EBS) volumes are encrypted at rest is to enable EBS encryption by default in all AWS Regions. This setting automatically encrypts all new EBS volumes and snapshots created in the account, thereby ensuring compliance with encryption policies without the need for manual intervention or additional monitoring.

AWS Documentation on Amazon EBS encryption provides guidance on enabling EBS encryption by default. This approach aligns with AWS best practices for data protection and compliance, ensuring that all new EBS volumes adhere to encryption requirements with minimal operational effort.

Question: 421

A company wants to establish a dedicated connection between its on-premises infrastructure and AWS. The company is setting up a 1 Gbps AWS Direct Connect connection to its account VPC. The architecture includes a transit gateway and a Direct Connect gateway to connect multiple VPCs and the on-premises infrastructure.

The company must connect to VPC resources over a transit VIF by using the Direct Connect connection.

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

- A. Update the 1 Gbps Direct Connect connection to 10 Gbps.
- B. Advertise the on-premises network prefixes over the transit VIF.
- C. Advertise the VPC prefixes from the Direct Connect gateway to the on-premises network over the transit VIF.
- D. Update the Direct Connect connection's MACsec encryption mode attribute to must encrypt.
- E. Associate a MACsec Connection Key Name-Connectivity Association Key (CKN/CAK) pair with the Direct Connect connection.

Answer: B,C

Explanation:

To connect VPC resources over a transit Virtual Interface (VIF) using a Direct Connect connection, the company should advertise the on-premises network prefixes over the transit VIF and advertise the VPC prefixes from the Direct Connect gateway to the on-premises network over the same VIF. This configuration ensures seamless connectivity between the on-premises infrastructure and the AWS VPCs through the transit gateway, facilitating efficient and secure communication across the network.

AWS Documentation on AWS Direct Connect and transit gateways provides detailed instructions on configuring transit VIFs and routing for Direct Connect connections. This setup is recommended in AWS best practices for establishing dedicated network connections between on-premises environments and AWS to achieve low-latency, high-throughput, and secure connectivity.

Question: 422

A company is designing an AWS environment for 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 will 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 Connect connection. Configure MACsec to encrypt traffic inside the Direct Connect connection.
- B. Provision another Direct Connect 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 Direct Connect 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:

To enhance the network connectivity solution's availability, fault tolerance, and security in a cost-effective manner, adding a dynamic private IP AWS Site-to-Site VPN as a secondary path is a viable option. This VPN serves as a resilient backup for the Direct Connect connection, ensuring continuous data flow even if the primary path fails. Implementing MACsec (Media Access Control Security) on the Direct Connect connection further secures the data in transit by providing encryption, thus addressing the security requirement. This solution strikes a balance between cost and operational efficiency, avoiding the higher expenses associated with provisioning an additional Direct Connect connection.

AWS Documentation on AWS Direct Connect and AWS Site-to-Site VPN provides insights into setting up resilient and secure network connections. Additionally, information on MACsec offers guidance on how to implement encryption for Direct Connect connections, aligning with best practices for secure and highly available network architectures.

Question: 423

A company needs to implement disaster recovery for a critical application that runs in a single AWS Region. The application's users interact with a web frontend that is hosted on Amazon EC2 Instances behind an Application Load Balancer (ALB). The application writes to an Amazon RDS for MySQL DB instance. The application also outputs processed documents that are stored in an Amazon S3 bucket. The company's finance team directly queries the database to run reports. During busy periods, these queries consume resources and negatively affect application performance.

A solutions architect must design a solution that will provide resiliency during a disaster. The solution must minimize data loss and must resolve the performance problems that result from the finance team's queries. Which solution will meet these requirements?

- A. Migrate the database to Amazon DynamoDB and use DynamoDB global tables. Instruct the finance team to query a global table in a separate Region. Create an AWS Lambda function to periodically synchronize the contents of the original S3 bucket to a new S3 bucket in the separate Region. Launch EC2 instances and create an ALB in the separate Region. Configure the application to point to the new S3 bucket.
- B. Launch additional EC2 instances that host the application in a separate Region. Add the additional instances to the existing ALB. In the separate Region, create a read replica of the RDS DB instance. Instruct the finance team to run queries against the read replica. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 bucket in the separate Region. During a disaster, promote the read replica to a standalone DB instance. Configure the application to point to the new S3 bucket and to the newly project read replica.
- C. Create a read replica of the RDS DB instance in a separate Region. Instruct the finance team to run queries against the read replica. Create AMIs of the EC2 instances that host the application frontend. Copy the AMIs to the separate Region. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 bucket in the separate Region. During a disaster, promote the read replica to a standalone DB instance. Launch EC2 instances from the AMIs and create an ALB to present the application to end users. Configure the application to point to the new S3 bucket.
- D. Create hourly snapshots of the RDS DB instance. Copy the snapshots to a separate Region. Add an Amazon ElastiCache cluster in front of the existing RDS database. Create AMIs of the EC2 instances that host the application frontend. Copy the AMIs to the separate Region. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 bucket in the separate Region. During a disaster, restore the database from the latest RDS snapshot. Launch EC2 instances from the AMIs and create an ALB to present the application to end users. Configure the application to point to the new S3 bucket.

Answer: C

Explanation:

Implementing a disaster recovery strategy that minimizes data loss and addresses performance issues involves creating a read replica of the RDS DB instance in a separate region and directing the finance team's queries to this replica. This solution alleviates the performance impact on the primary database. Using Amazon S3 Cross-Region Replication (CRR) ensures that processed documents are available in the disaster recovery region. In the event of a disaster, the read replica can be promoted to a standalone DB instance, and EC2 instances can be launched from pre-created AMIs to serve the web frontend, thereby ensuring resiliency and minimal data loss.

AWS Documentation on Amazon RDS Read Replicas, Amazon S3 Cross-Region Replication, and Amazon EC2

AMIs provides comprehensive guidance on implementing a robust disaster recovery solution. This approach is in line with AWS best practices for high availability and disaster recovery planning.

Question: 424

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 of 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:

Utilizing an IP access control group rule with the list of public addresses from branch offices and associating it with the Amazon WorkSpaces directory is the most operationally efficient solution. This method ensures that access to WorkSpaces is restricted to specified locations, aligning with the corporate security policy. This approach offers simplicity and flexibility, especially with the potential addition of a new branch office, as updating the IP access control group is straightforward.

AWS Documentation on Amazon WorkSpaces and IP Access Control Groups provides detailed instructions on how to implement access restrictions based on IP addresses. This solution aligns with best practices for securing virtual desktops while maintaining operational efficiency.

Question: 425

A company has a Windows-based desktop application that is packaged and deployed to the users' Windows machines. The company recently acquired another company that has employees who primarily use machines with a Linux operating system. The acquiring company has decided to migrate and rehost the Windows-based desktop application to AWS.

All employees must be authenticated before they use the application. The acquiring company uses Active Directory on premises but wants a simplified way to manage access to the application on AWS (or all the employees).

Which solution will rehost the application on AWS with the LEAST development effort?

- A. Set up and provision an Amazon Workspaces virtual desktop for every employee. Implement authentication

by using Amazon Cognito identity pools. Instruct employees to run the application from their provisioned Workspaces virtual desktops.

B. Create an Auto Scaling group of Windows-based Amazon EC2 instances. Join each EC2 instance to the company's Active Directory domain. Implement authentication by using the Active Directory That is running on premises. Instruct employees to run the application by using a Windows remote desktop.

C. Use an Amazon AppStream 2.0 image builder to create an image that includes the application and the required configurations. Provision an AppStream 2.0 On-Demand fleet with dynamic Fleet Auto Scaling process for running the image. Implement authentication by using AppStream 2.0 user pools. Instruct the employees to access the application by starting browser-based AppStream 2.0 streaming sessions.

D. Refactor and containerize the application to run as a web-based application. Run the application in Amazon Elastic Container Service (Amazon ECS) on AWS Fargate with step scaling policies. Implement authentication by using Amazon Cognito user pools. Instruct the employees to run the application from their browsers.

Answer: C

Explanation:

Amazon AppStream 2.0 offers a streamlined solution for rehosting a Windows-based desktop application on AWS with minimal development effort. By creating an AppStream 2.0 image that includes the application and using an On-Demand fleet for streaming, the application becomes accessible from any device, including Linux machines. AppStream 2.0 user pools can be used for authentication, simplifying access management without the need for extensive changes to the application or infrastructure.

AWS Documentation on Amazon AppStream 2.0 provides insights into setting up application streaming solutions. This approach is recommended for delivering desktop applications to diverse operating systems without the complexity of managing virtual desktops or extensive application refactoring.

Question: 426

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

Which solution will meet these requirements?

A. Order an AWS Snowball Edge Compute Optimized device. Connect the device to the local network. Configure AWS DataSync with a target bucket name, and upload the data over NFS to the device. After the experiment return the device to AWS so that the data can be loaded into Amazon S3.

B. Order an AWS Snowcone device, including an Amazon Linux 2 AMI. Connect the device to the local network. Launch an Amazon EC2 instance on the device. Create a shell script that periodically downloads data from each sensor. After the experiment, return the device to AWS so that the data can be loaded as an Amazon Elastic Block Store [Amazon EBS] volume.

C. Order an AWS Snowcone device, including an Amazon Linux 2 AMI. Connect the device to the local network. Launch an Amazon EC2 instance on the device. Install and configure an FTP server on the EC2 instance. Configure the sensors to upload data to the EC2 instance. After the experiment, return the device to AWS

so that the data can be loaded into Amazon S3.

D. Order an AWS Snowcone device. Connect the device to the local network. Configure the device to use Amazon FSx. Configure the sensors to upload data to the device. Configure AWS DataSync on the device to synchronize the uploaded data with an Amazon S3 bucket. Return the device to AWS so that the data can be loaded as an Amazon Elastic Block Store (Amazon EBS) volume.

Answer: C

Explanation:

For collecting data from remote sensors without internet connectivity, using an AWS Snowcone device with an Amazon EC2 instance running an FTP server presents a practical solution. This setup allows the sensors to upload data to the EC2 instance via FTP, and after the experiment, the Snowcone device can be returned to AWS for data ingestion into Amazon S3. This approach minimizes operational complexity and ensures efficient data transfer to AWS for further processing or storage.

AWS Documentation on AWS Snowcone and Amazon EC2 provides detailed guidance on deploying compute and storage capabilities in edge locations. This solution leverages AWS's edge computing devices to address challenges associated with data collection in remote or disconnected environments.

Question: 427

A flood monitoring agency has deployed more than 10,000 water-level monitoring sensors. Sensors send continuous data updates, and each update is less than 1 MB in size. The agency has a fleet of on-premises application servers. These servers receive updates from the sensors, convert the raw data into a human readable format, and write the results to an on-premises relational database server. Data analysts then use simple SQL queries to monitor the data.

The agency wants to increase overall application availability and reduce the effort that is required to perform maintenance tasks. These maintenance tasks, which include updates and patches to the application servers, cause downtime. While an application server is down, data is lost from sensors because the remaining servers cannot handle the entire workload.

The agency wants a solution that optimizes operational overhead and costs. A solutions architect recommends the use of AWS IoT Core to collect the sensor data.

What else should the solutions architect recommend to meet these requirements?

A. Send the sensor data to Amazon Kinesis Data Firehose. Use an AWS Lambda function to read the Kinesis Data Firehose data, convert it to .csv format, and insert it into an Amazon Aurora MySQL DB instance. Instruct the data analysts to query the data directly from the DB instance.

B. Send the sensor data to Amazon Kinesis Data Firehose. Use an AWS Lambda function to read the Kinesis Data Firehose data, convert it to Apache Parquet format and save it to an Amazon S3 bucket. Instruct the data analysts to query the data by using Amazon Athena.

C. Send the sensor data to an Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) application to convert the data to .csv format and store it in an Amazon S3 bucket. Import the data into an Amazon Aurora MySQL DB instance. Instruct the data analysts to query the data directly from the DB instance.

D. Send the sensor data to an Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) application to convert the data to Apache Parquet format and store it in an Amazon S3 bucket.

Instruct the data analysis to query the data by using Amazon Athena.

Answer: B

Explanation:

To enhance application availability and reduce maintenance-induced downtime, sending sensor data to Amazon Kinesis Data Firehose, processing it with an AWS Lambda function, converting it to Apache Parquet format, and storing it in Amazon S3 is an effective strategy. This approach leverages serverless architectures for scalability and reliability. Data analysts can then query the optimized data using Amazon Athena, a serverless interactive query service, which supports complex queries on data stored in S3 without the need for traditional database servers, optimizing operational overhead and costs.

AWS Documentation on AWS IoT Core, Amazon Kinesis Data Firehose, AWS Lambda, Amazon S3, and Amazon Athena provides a comprehensive framework for building a scalable, serverless data processing pipeline. This solution aligns with AWS best practices for processing and analyzing large-scale data streams efficiently.

Question: 428

A company needs to improve the reliability ticketing application. The application runs on an Amazon Elastic Container Service (Amazon ECS) cluster. The company uses Amazon CloudFront to serve the application. A single ECS service of the ECS cluster is the CloudFront distribution's origin.

The application allows only a specific number of active users to enter a ticket purchasing flow. These users are identified by an encrypted attribute in their JSON Web Token (JWT). All other users are redirected to a waiting room module until there is available capacity for purchasing.

The application is experiencing high loads. The waiting room module is working as designed, but load on the waiting room is disrupting the application's availability. This disruption is negatively affecting the application's ticket sale Transactions.

Which solution will provide the MOST reliability for ticket sale transactions during periods of high load? '

- A. Create a separate service in the ECS cluster for the waiting room. Use a separate scaling configuration. Ensure that the ticketing service uses the JWT information and appropriately forwards requests to the waiting room service.
- B. Move the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Split the waiting room module into a pod that is separate from the ticketing pod. Make the ticketing pod part of a StatefulSet. Ensure that the ticketing pod uses the JWT information and appropriately forwards requests to the waiting room pod.
- C. Create a separate service in the ECS cluster for the waiting room. Use a separate scaling configuration. Create a CloudFront function that inspects the JWT information and appropriately forwards requests to the ticketing service or the waiting room service.
- D. Move the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Split the waiting room module into a pod that is separate from the ticketing pod. Use AWS App Mesh by provisioning the App Mesh controller for Kubernetes. Enable mTLS authentication and service-to-service authentication for communication between the ticketing pod and the waiting room pod. Ensure that the ticketing pod uses the JWT information and appropriately forwards requests to the waiting room pod.

Answer: C

Explanation:

Implementing a CloudFront function that inspects the JWT information and appropriately forwards requests either to the ticketing service or the waiting room service within the Amazon ECS cluster enhances reliability during high load periods. This solution segregates the load between the main application and the waiting room, ensuring that the ticketing service remains unaffected by the high load on the waiting room. Using CloudFront functions for request routing based on JWT attributes allows for efficient distribution of user traffic, thereby maintaining the application's availability and performance during peak times.

AWS Documentation on Amazon CloudFront Functions provides guidance on creating and deploying functions that can inspect and manipulate HTTP(S) requests at the edge, close to the users. This approach is in line with best practices for scaling and managing high-traffic web applications.

Question: 429

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

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

Answer: B

Explanation:

Setting up an AWS Client VPN endpoint and integrating it with Active Directory Domain Services (AD DS) using an AD Connector directory enables secure remote access to internal services deployed in a VPC. Enabling multi-factor authentication (MFA) for AD Connector enhances security by adding an additional layer of authentication. This solution meets the company's requirements for secure remote access through a VPN with MFA, ensuring that the security policy is adhered to while providing a seamless experience for the remote engineers.

AWS Documentation on AWS Client VPN and AD Connector provides detailed instructions on setting up a Client VPN endpoint and integrating it with existing Active Directory for authentication. This solution aligns with AWS best practices for secure remote access to AWS resources.

Question: 430

A company deploys workloads in multiple AWS accounts. Each account has a VPC with VPC flow logs published in text log format to a centralized Amazon S3 bucket. Each log file is compressed with gzip compression. The company must retain the log files indefinitely.

A security engineer occasionally analyzes the logs by using Amazon Athena to query the VPC flow logs. The query performance is degrading over time as the number of ingested logs is growing. A solutions architect must improve the performance of the log analysis and reduce the storage space that the VPC flow logs use.

Which solution will meet these requirements with the LARGEST performance improvement?

- A. Create an AWS Lambda function to decompress the gzip files and to compress the files with bzip2 compression. Subscribe the Lambda function to an s3:ObjectCreated:Put S3 event notification for the S3 bucket.
- B. Enable S3 Transfer Acceleration for the S3 bucket. Create an S3 Lifecycle configuration to move files to the S3 Intelligent-Tiering storage class as soon as the files are uploaded.
- C. Update the VPC flow log configuration to store the files in Apache Parquet format. Specify Hourly partitions for the log files.
- D. Create a new Athena workgroup without data usage control limits. Use Athena engine version 2.

Answer: C

Explanation:

Converting VPC flow logs to store in Apache Parquet format and specifying hourly partitions significantly improves query performance and reduces storage space usage. Apache Parquet is a columnar storage file format optimized for analytical queries, allowing Athena to scan less data and improve query performance. Partitioning logs by hour further enhances query efficiency by limiting the amount of data scanned during queries, addressing the issue of degrading performance over time due to the growing volume of ingested logs.

AWS Documentation on VPC Flow Logs and Amazon Athena provides insights into configuring VPC flow logs in Apache Parquet format and using Athena for querying log data. This approach is recommended for efficient log analysis and storage optimization.

Question: 431

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

Auto Scaling is configured to maintain the web fleet size based on the ALB health check.

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

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

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

Answer: A,E

Explanation:

Configuring read replicas for Amazon RDS MySQL and using the single reader endpoint in the web application can significantly reduce the load on the backend database tier, improving overall application performance.

Additionally, implementing an Amazon ElastiCache cluster between the web application and RDS MySQL instances can further reduce database load by caching frequently accessed data, thereby enhancing the application's resilience and scalability. These changes address the root cause of the outage by alleviating the database tier's high load and preventing similar issues in the future.

AWS Documentation on Amazon RDS Read Replicas and Amazon ElastiCache provides comprehensive guidance on improving application performance and scalability by offloading read traffic from the primary database and caching common queries. These solutions are in line with AWS best practices for building resilient and scalable web applications.

Question: 432

A company is running its solution on AWS in a manually created VPC. The company is using AWS

CloudFormation to provision other parts of the infrastructure. According to a new requirement, the company must manage all infrastructure in an automatic way.

What should the company do to meet this new requirement with the LEAST effort?

- A. Create a new AWS Cloud Development Kit (AWS CDK) stack that strictly provisions the existing VPC resources and configuration. Use AWS CDK to import the VPC into the stack and to manage the VPC.
- B. Create a CloudFormation stack set that creates the VPC. Use the stack set to import the VPC into the stack.
- C. Create a new CloudFormation template that strictly provisions the existing VPC resources and configuration. From the CloudFormation console, create a new stack by importing the existing resources.
- D. Create a new CloudFormation template that creates the VPC. Use the AWS Serverless Application Model (AWS SAM) CLI to import the VPC.

Answer: C

Explanation:

Creating the Template:

Start by creating a CloudFormation template that includes all the VPC resources. This template should accurately reflect the current state and configuration of the VPC.

Using the CloudFormation Console:

Open the AWS Management Console and navigate to CloudFormation.

Choose "Create stack" and then select "With existing resources (import resources)".

Specifying the Template:

Upload the previously created template or specify the Amazon S3 URL where the template is stored.

Identifying the Resources:

On the "Identify resources" page, provide the identifiers for each VPC resource you wish to import. For example, for an AWS::EC2::VPC resource, use the VPC ID as the identifier.

Creating the Stack:

Complete the stack creation process by providing stack details and reviewing the changes. This will create a change set that includes the import operation.

Executing the Change Set:

Execute the change set to import the resources into the CloudFormation stack, making them managed by CloudFormation.

Verification and Drift Detection:

After the import is complete, run drift detection to ensure the actual configuration matches the template configuration.

This approach allows the company to manage their VPC and other resources via CloudFormation without the need to recreate resources, ensuring a smooth transition to automated infrastructure management.

Reference

Creating a stack from existing resources - AWS CloudFormation(AWS Documentation). Generating templates for existing resources - AWS CloudFormation(AWS Documentation).

Bringing existing resources into CloudFormation management(AWS Documentation).

Question: 433

A company runs a software-as-a-service (SaaS) application on AWS. The application consists of AWS Lambda functions and an Amazon RDS for MySQL Multi-AZ database. During market events, the application has a much higher workload than normal. Users notice slow response times during the peak periods because of many database connections. The company needs to improve the scalable performance and availability of the database.

Which solution meets these requirements?

- A. Create an Amazon CloudWatch alarm action that triggers a Lambda function to add an Amazon RDS for MySQL read replica when resource utilization hits a threshold.
- B. Migrate the database to Amazon Aurora, and add a read replica. Add a database connection pool outside of the Lambda handler function.
- C. Migrate the database to Amazon Aurora and add a read replica. Use Amazon Route 53 weighted records.
- D. Migrate the database to Amazon Aurora and add an Aurora Replica. Configure Amazon RDS Proxy to manage database connection pools.

Answer: D

Explanation:

Migrate to Amazon Aurora:

Amazon Aurora is a MySQL-compatible, high-performance database designed to provide higher throughput than standard MySQL. Migrating the database to Aurora will enhance the performance and scalability of the database, especially under heavy workloads.

Add Aurora Replica:

Aurora Replicas provide read scalability and improve availability. Adding an Aurora Replica allows read operations to be distributed, thereby reducing the load on the primary instance and improving response times during peak periods.

Configure Amazon RDS Proxy:

Amazon RDS Proxy acts as an intermediary between the application and the Aurora database, managing connection pools efficiently. RDS Proxy reduces the overhead of opening and closing database connections, thus maintaining fewer active connections to the database and handling surges in database connections from the Lambda functions more effectively.

This configuration reduces the database's resource usage and improves its ability to handle high volumes of concurrent connections.

Reference

AWS Database Blog on RDS Proxy(Amazon Web Services, Inc.).

AWS Compute Blog on Using RDS Proxy with Lambda(Amazon Web Services, Inc.).

Question: 434

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

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

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

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

Answer: A,C,F

Explanation:

Auto Scaling Group with Application Load Balancer:

Moving the Tomcat server to an Auto Scaling group ensures that the number of instances adjusts dynamically based on the traffic load. An Application Load Balancer (ALB) distributes incoming traffic across multiple instances, improving the application's reliability and availability.

Migrate to Amazon Aurora with Replica:

Migrating the MySQL database to Amazon Aurora and adding an Aurora Replica enhances the database's scalability and availability. Aurora is optimized for performance, and replicas help distribute read traffic, reducing the load on the primary instance.

Additional Public Subnet:

Creating an additional public subnet in a different Availability Zone enhances fault tolerance. This ensures that the website remains accessible even if one Availability Zone experiences issues. [Reference](#)

[AWS Well-Architected Framework](#)

[Amazon Aurora Documentation](#)

Question: 435

A company wants to migrate virtual Microsoft workloads from an on-premises data center to AWS. The company has successfully tested a few sample workloads on AWS. The company also has created an AWS Site-to-Site VPN connection to a VPC. A solutions architect needs to generate a total cost of ownership (TCO) report for the migration of all the workloads from the data center.

Simple Network Management Protocol (SNMP) has been enabled on each VM in the data center. The company cannot add more VMs in the data center and cannot install additional software on the VMs. The discovery data must be automatically imported into AWS Migration Hub.

Which solution will meet these requirements?

- A. Use the AWS Application Migration Service agentless service and the AWS Migration Hub Strategy Recommendations to generate the TCO report.
- B. Launch a Windows Amazon EC2 instance. Install the Migration Evaluator agentless collector on the EC2 instance. Configure Migration Evaluator to generate the TCO report.
- C. Launch a Windows Amazon EC2 instance. Install the Migration Evaluator agentless collector on the EC2 instance. Configure Migration Hub to generate the TCO report.
- D. Use the AWS Migration Readiness Assessment tool inside the VPC. Configure Migration Evaluator to generate the TCO report.

Answer: A

Explanation:

AWS Application Migration Service:

AWS Application Migration Service (MGN) facilitates the migration of virtual machines (VMs) to AWS without installing additional software on the VMs. This agentless service helps in seamlessly migrating workloads to AWS.

AWS Migration Hub Strategy Recommendations:

AWS Migration Hub Strategy Recommendations offer insights and guidance for planning and implementing migration strategies. It helps in generating a Total Cost of Ownership (TCO) report by automatically importing discovery data from the VMs.

Generating the TCO Report:

The combined use of AWS Application Migration Service and Migration Hub Strategy Recommendations enables the automatic import of discovery data and the generation of an accurate TCO report, ensuring a smooth and cost-effective migration process.

Reference

AWS Migration Hub Strategy Recommendations(AWS Documentation).

Question: 436

A startup company recently migrated a large ecommerce website to AWS The website has experienced a 70% increase in sales Software engineers are using a private GitHub repository to manage code The DevOps team is using Jenkins for builds and unit testing The engineers need to receive notifications for bad builds and zero downtime during deployments The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue

The software engineers have decided to use AWS CodePipeline to manage their build and deployment process

Which solution will meet these requirements'?

- A. Use GitHub websockets to trigger the CodePipeline pipeline Use the Jenkins plugin for AWS CodeBuild to conduct unit testing Send alerts to an Amazon SNS topic for any bad builds Deploy in an in-place all-at-once deployment configuration using AWS CodeDeploy
- B. Use GitHub webhooks to trigger the CodePipeline pipeline Use the Jenkins plugin for AWS CodeBuild to conduct unit testing Send alerts to an Amazon SNS topic for any bad builds Deploy in a blue/green deployment using AWS CodeDeploy
- C. Use GitHub websockets to trigger the CodePipeline pipeline. Use AWS X-Ray for unit testing and static code analysis Send alerts to an Amazon SNS topic for any bad builds Deploy in a blue/green deployment using AWS CodeDeploy.
- D. Use GitHub webhooks to trigger the CodePipeline pipeline Use AWS X-Ray for unit testing and static code analysis Send alerts to an Amazon SNS topic for any bad builds Deploy in an m-place. all- at-once deployment configuration using AWS CodeDeploy

Answer: B

Explanation:

GitHub Webhooks to Trigger CodePipeline:

Configure GitHub webhooks to trigger the AWS CodePipeline pipeline. This ensures that every code push to the repository automatically triggers the pipeline, initiating the build and deployment process.

Unit Testing with Jenkins and AWS CodeBuild:

Use Jenkins integrated with the AWS CodeBuild plugin to perform unit testing. Jenkins will manage the build process, and the results will be handled by CodeBuild.

Notifications for Bad Builds:

Configure Amazon SNS (Simple Notification Service) to send alerts for any failed builds. This keeps the engineering team informed of build issues immediately, allowing for quick resolutions. Blue/Green Deployment with AWS CodeDeploy:

Utilize AWS CodeDeploy with a blue/green deployment strategy. This method reduces downtime and risk by running two identical production environments (blue and green). CodeDeploy shifts traffic between these

environments, allowing you to test in the new environment (green) while the old environment (blue) remains live. If issues arise, you can quickly roll back to the previous environment.

This solution provides seamless, zero-downtime deployments, and the ability to quickly roll back changes if necessary, fulfilling the requirements of the startup company.

Reference

AWS DevOps Blog on Integrating Jenkins with AWS CodeBuild and CodeDeploy [32] .

Plain English Guide to AWS CodePipeline with GitHub [33] .

Jenkins Plugin for AWS CodePipeline[34] .

Question: 437

A software as a service (SaaS) company has developed a multi-tenant environment. The company uses Amazon DynamoDB tables that the tenants share for the storage layer. The company uses AWS Lambda functions for the application services.

The company wants to offer a tiered subscription model that is based on resource consumption by each tenant. Each tenant is identified by a unique tenant ID that is sent as part of each request to the Lambda functions. The company has created an AWS Cost and Usage Report (AWS CUR) in an AWS account. The company wants to allocate the DynamoDB costs to each tenant to match that tenant's resource consumption.

Which solution will provide a granular view of the DynamoDB cost for each tenant with the LEAST operational effort?

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

Answer: B

Explanation:

Log Tenant ID and RCUs/WCUs:

Update the AWS Lambda functions to log the tenant ID and the number of Read Capacity Units (RCUs) and

Write Capacity Units (WCUs) consumed from DynamoDB for each transaction. This data will be logged to Amazon CloudWatch Logs.

Calculate Tenant Costs:

Deploy an additional Lambda function that reads the logs from CloudWatch Logs, calculates the RCUs and WCUs used by each tenant, and then uses the AWS Cost Explorer API to retrieve the overall cost of DynamoDB usage.

This function will then allocate the costs to each tenant based on their usage. Scheduled Cost Calculation:

Create an Amazon EventBridge rule to trigger the cost calculation Lambda function at regular intervals (e.g., daily or hourly). This ensures that cost allocation is continuously updated and tenants are billed accurately based on their consumption.

This solution minimizes operational effort by automating the cost allocation process and ensuring that the company can accurately bill tenants based on their resource consumption.

Reference

[AWS Cost Explorer Documentation](#)

[Amazon CloudWatch Logs Documentation](#)

[AWS Lambda Documentation](#)

Question: 438

A company that develops consumer electronics with offices in Europe and Asia has 60 TB of software images stored on premises in Europe. The company wants to transfer the images to an Amazon S3 bucket in the ap-northeast-1 Region. New software images are created daily and must be encrypted in transit. The company needs a solution that does not require custom development to automatically transfer all existing and new software images to Amazon S3. What is the next step in the transfer process?

- A. Deploy an AWS DataSync agent and configure a task to transfer the images to the S3 bucket
- B. Configure Amazon Kinesis Data Firehose to transfer the images using S3 Transfer Acceleration
- C. Use an AWS Snowball device to transfer the images with the S3 bucket as the target
- D. Transfer the images over a Site-to-Site VPN connection using the S3 API with multipart upload

Answer: A

Explanation:

Deploy AWS DataSync Agent:

Install the DataSync agent on your on-premises environment. This can be done by downloading the agent as a virtual appliance and deploying it on VMware ESXi, Hyper-V, or KVM hypervisors. Configure Source and

Destination Locations:

Set up the source location pointing to your on-premises storage where the software images are currently stored.

Configure the destination location to point to your Amazon S3 bucket in the ap-northeast-1 Region. Create and Schedule DataSync Tasks:

Create a DataSync task to automate the transfer process. This task will specify the source and destination locations and set options for how the data should be transferred.

Schedule the task to run at intervals that suit your data transfer requirements, ensuring new images are transferred as they are created.

Encryption in Transit:

AWS DataSync automatically encrypts data in transit using TLS, ensuring that your data is secure during the transfer process.

Monitoring and Management:

Use the DataSync console or the AWS CLI to monitor the progress of your data transfers and manage the tasks.

AWS DataSync is an efficient solution that automates and accelerates the process of transferring large amounts of data to AWS, handling encryption, data integrity checks, and optimizing network usage without requiring custom development.

Reference

AWS Storage Blog on DataSync [40] .

AWS DataSync Documentation [41] .

Question: 439

An events company runs a ticketing platform on AWS. The company's customers configure and schedule their events on the platform. The events result in large increases of traffic to the platform. The company knows the date and time of each customer's events.

The company runs the platform on an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS cluster consists of Amazon EC2 On-Demand Instances that are in an Auto Scaling group. The Auto Scaling group uses a predictive scaling policy.

The ECS cluster makes frequent requests to an Amazon S3 bucket to download ticket assets. The ECS cluster and the S3 bucket are in the same AWS Region and the same AWS account. Traffic between the ECS cluster and the S3 bucket flows across a NAT gateway.

The company needs to optimize the cost of the platform without decreasing the platform's availability. Which combination of steps will meet these requirements? (Select TWO)

- A. Create a gateway VPC endpoint for the S3 bucket.
- B. Add another ECS capacity provider that uses an Auto Scaling group of Spot Instances. Configure the new capacity provider strategy to have the same weight as the existing capacity provider strategy.
- C. Create On-Demand Capacity Reservations for the applicable instance type for the time period of the scheduled scaling policies.
- D. Enable S3 Transfer Acceleration on the S3 bucket.
- E. Replace the predictive scaling policy with scheduled scaling policies for the scheduled events.

Answer: A,B

Explanation:

Gateway VPC Endpoint for S3:

Create a gateway VPC endpoint for Amazon S3 in your VPC. This allows instances in your VPC to communicate with Amazon S3 without going through the internet, reducing data transfer costs and improving security.

Add Spot Instances to ECS Cluster:

Add another ECS capacity provider that uses an Auto Scaling group of Spot Instances. Configure this new capacity provider to share the load with the existing On-Demand Instances by setting an appropriate weight in the capacity provider strategy. Spot Instances offer significant cost savings compared to On-Demand

Instances.

Configure Capacity Provider Strategy:

Adjust the ECS service's capacity provider strategy to utilize both On-Demand and Spot Instances effectively. This ensures a balanced distribution of tasks across both instance types, optimizing cost while maintaining availability.

By implementing a gateway VPC endpoint for S3 and incorporating Spot Instances into the ECS cluster, the company can significantly reduce operational costs without compromising on the availability or performance of the platform.

Reference

AWS Cost Optimization Blog on VPC Endpoints

AWS ECS Documentation on Capacity Providers

Question: 440

A company uses AWS Organizations to manage its development environment. Each development team at the company has its own AWS account. Each account has a single VPC and CIDR blocks that do not overlap. The company has an Amazon Aurora DB cluster in a shared services account. All the development teams need to work with live data from the DB cluster.

Which solution will provide the required connectivity to the DB cluster with the LEAST operational overhead?

- A. Create an AWS Resource Access Manager (AWS RAM) resource share for the DB cluster. Share the DB cluster with all the development accounts.
- B. Create a transit gateway in the shared services account. Create an AWS Resource Access Manager (AWS RAM) resource share for the transit gateway. Share the transit gateway with all the development accounts. Instruct the developers to accept the resource share. Configure networking.
- C. Create an Application Load Balancer (ALB) that points to the IP address of the DB cluster. Create an AWS PrivateLink endpoint service that uses the ALB. Add permissions to allow each development account to connect to the endpoint service.
- D. Create an AWS Site-to-Site VPN connection in the shared services account. Configure networking. Use AWS Marketplace VPN software in each development account to connect to the Site-to-Site VPN connection.

Answer: B

Explanation:

Create a Transit Gateway:

In the shared services account, create a new AWS Transit Gateway. This serves as a central hub to connect multiple VPCs, simplifying the network topology and management.

Configure Transit Gateway Attachments:

Attach the VPC containing the Aurora DB cluster to the transit gateway. This allows the shared services VPC to communicate through the transit gateway.

Create Resource Share with AWS RAM:

Use AWS Resource Access Manager (AWS RAM) to create a resource share for the transit gateway. Share this resource with all development accounts. AWS RAM allows you to securely share your AWS resources across AWS accounts without needing to duplicate them.

Accept Resource Shares in Development Accounts:

Instruct each development team to log into their respective AWS accounts and accept the transit gateway.

resource share. This step is crucial for enabling cross-account access to the shared transit gateway.

Configure VPC Attachments in Development Accounts:

Each development account needs to attach their VPC to the shared transit gateway. This allows their VPCs to route traffic through the transit gateway to the Aurora DB cluster in the shared services account.

Update Route Tables:

Update the route tables in each VPC to direct traffic intended for the Aurora DB cluster through the transit gateway. This ensures that network traffic is properly routed between the development VPCs and the shared services VPC.

Using a transit gateway simplifies the network management and reduces operational overhead by providing a scalable and efficient way to interconnect multiple VPCs across different AWS accounts. Reference

AWS Database Blog on RDS Proxy for Cross-Account Access [48] .

AWS Architecture Blog on Cross-Account and Cross-Region Aurora Setup [49] .

DEV Community on Managing Multiple AWS Accounts with Organizations [51] .

Question: 441

A company has an application that analyzes and stores image data on premises. The application receives millions of new image files every day. Files are an average of 1 MB in size. The files are analyzed in batches of 1 GB. When the application analyzes a batch, the application zips the images together. The application then archives the images as a single file in an on-premises NFS server for long-term storage.

The company has a Microsoft Hyper-V environment on premises and has compute capacity available. The company does not have storage capacity and wants to archive the images on AWS. The company needs the ability to retrieve archived data within a week of a request.

The company has a 10 Gbps AWS Direct Connect connection between its on-premises data center and AWS. The company needs to set bandwidth limits and schedule archived images to be copied to AWS during non-business hours.

Which solution will meet these requirements MOST cost-effectively?

A. Deploy an AWS DataSync agent on a new GPU-based Amazon EC2 instance. Configure the DataSync agent to copy the batch of files from the NFS on-premises server to Amazon S3 Glacier Instant Retrieval. After the successful copy, delete the data from the on-premises storage.

B. Deploy an AWS DataSync agent as a Hyper-V VM on premises. Configure the DataSync agent to copy the batch of files from the NFS on-premises server to Amazon S3 Glacier Deep Archive. After the successful copy, delete the data from the on-premises storage.

C. Deploy an AWS DataSync agent on a new general purpose Amazon EC2 instance. Configure the DataSync agent to copy the batch of files from the NFS on-premises server to Amazon S3 Standard. After the successful copy, delete the data from the on-premises storage. Create an S3 Lifecycle rule to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 day.

D. Deploy an AWS Storage Gateway Tape Gateway on premises in the Hyper-V environment. Connect the Tape Gateway to AWS. Use automatic tape creation. Specify an Amazon S3 Glacier Deep Archive pool. Eject the tape after the batch of images is copied.

Answer: B

Explanation:

Deploy DataSync Agent:

Install the AWS DataSync agent as a VM in your Hyper-V environment. This agent facilitates the data transfer between your on-premises storage and AWS.

Configure Source and Destination:

Set up the source location to point to your on-premises NFS server where the image batches are stored. Configure the destination location to be an Amazon S3 bucket with the Glacier Deep Archive storage class. This storage class is cost-effective for long-term storage with retrieval times of up to 12 hours. **Create DataSync Tasks:**

Create and configure DataSync tasks to manage the data transfer. Schedule these tasks to run during non-business hours to minimize bandwidth usage during peak times. The tasks will handle the copying of data batches from the NFS server to the S3 bucket.

Set Bandwidth Limits:

In the DataSync configuration, set bandwidth limits to control the amount of data being transferred at any given time. This ensures that your network's performance is not adversely affected during business hours.

Delete On-Premises Data:

After successfully copying the data to S3 Glacier Deep Archive, configure the DataSync task to delete the data from your on-premises NFS server. This helps manage storage capacity on-premises and ensures data is securely archived on AWS.

This approach leverages AWS DataSync for efficient, secure, and automated data transfer, and S3 Glacier Deep Archive for cost-effective long-term storage.

Reference

AWS DataSync Overview [41] .

AWS Storage Blog on DataSync Migration [40] .

Amazon S3 Transfer Acceleration Documentation [42] .

Question: 442

A company has a web application that uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. A recent marketing campaign has increased demand. Monitoring software reports that many requests have significantly longer response times than before the marketing campaign. A solutions architect enabled Amazon CloudWatch Logs for API Gateway and noticed that errors are occurring on 20% of the requests. In CloudWatch, the Lambda function's Throttles metric represents 1% of the requests and the Errors metric represents 10% of the requests. Application logs indicate that, when errors occur, there is a call to DynamoDB. What change should the solutions architect make to improve the current response times as the web application becomes more popular?

- A. Increase the concurrency limit of the Lambda function.
- B. Implement DynamoDB auto scaling on the table.
- C. Increase the API Gateway throttle limit.
- D. Re-create the DynamoDB table with a better-partitioned primary index.

Answer: B

Explanation:

Enable DynamoDB Auto Scaling:

Navigate to the DynamoDB console and select the table experiencing high demand.

Go to the "Capacity" tab and enable auto scaling for both read and write capacity units. Auto scaling adjusts the provisioned throughput capacity automatically in response to actual traffic patterns, ensuring the table can handle the increased load.

Configure Auto Scaling Policies:

Set the minimum and maximum capacity units to define the range within which auto scaling can adjust the provisioned throughput.

Specify target utilization percentages for read and write operations, typically around 70%, to maintain a balance between performance and cost.

Monitor and Adjust:

Use Amazon CloudWatch to monitor the auto scaling activity and ensure it is effectively handling the increased demand.

Adjust the auto scaling settings if necessary to better match the traffic patterns and application requirements.

By enabling DynamoDB auto scaling, you ensure that the database can handle the fluctuating traffic volumes without manual intervention, improving response times and reducing errors.

Reference

[AWS Compute Blog on Using API Gateway as a Proxy for DynamoDB \[60\]](#) .

[AWS Database Blog on DynamoDB Accelerator \(DAX\) \[59\]](#) .

Question: 443

A solutions architect has deployed a web application that serves users across two AWS Regions under a custom domain. The application uses Amazon Route 53 latency-based routing. The solutions architect has associated weighted record sets with a pair of web servers in separate Availability Zones for each Region.

The solutions architect runs a disaster recovery scenario. When all the web servers in one Region are stopped, Route 53 does not automatically redirect users to the other Region. Which of the following are possible root causes of this issue? (Select TWO)

- A. The weight for the Region where the web servers were stopped is higher than the weight for the other Region.
- B. One of the web servers in the secondary Region did not pass its HTTP health check.
- C. Latency resource record sets cannot be used in combination with weighted resource record sets.
- D. The setting to evaluate target health is not turned on for the latency alias resource record set that is associated with the domain in the Region where the web servers were stopped.
- E. An HTTP health check has not been set up for one or more of the weighted resource record sets associated with the stopped web servers.

Answer: D,E

Explanation:

Evaluate Target Health Setting:

Ensure that the "Evaluate Target Health" setting is enabled for the latency alias resource record sets in Route

53. This setting helps Route 53 determine the health of the resources associated with the alias record and redirect traffic appropriately.

HTTP Health Checks:

Configure HTTP health checks for all weighted resource record sets. Health checks monitor the availability and performance of the web servers, allowing Route 53 to reroute traffic to healthy servers in case of a failure. Verify that the health checks are correctly set up and associated with the resource record sets. This ensures that Route 53 can detect server failures and redirect traffic to the servers in the other Region.

By enabling the "Evaluate Target Health" setting and configuring HTTP health checks, Route 53 can effectively manage traffic during failover scenarios, ensuring high availability and reliability. Reference

AWS Route 53 Documentation on Latency-Based Routing[50] .

AWS Architecture Blog on Cross-Account and Cross-Region Setup[49] .

Question: 444

A solutions architect is creating an AWS CloudFormation template from an existing manually created non-production AWS environment. The CloudFormation template can be destroyed and recreated as needed. The environment contains an Amazon EC2 instance. The EC2 instance has an instance profile that the EC2 instance uses to assume a role in a parent account.

The solutions architect recreates the role in a CloudFormation template and uses the same role name. When the CloudFormation template is launched in the child account, the EC2 instance can no longer assume the role in the parent account because of insufficient permissions.

What should the solutions architect do to resolve this issue?

- A. In the parent account edit the trust policy for the role that the EC2 instance needs to assume. Ensure that the target role ARN in the existing statement that allows the sts AssumeRole action is correct. Save the trust policy.
- B. In the parent account edit the trust policy for the role that the EC2 instance needs to assume. Add a statement that allows the sts AssumeRole action for the root principal of the child account. Save the trust policy.
- C. Update the CloudFormation stack again. Specify only the CAPABILITY_NAMED_IAM capability.
- D. Update the CloudFormation stack again. Specify the CAPABILITY_IAM capability and the CAPABILITY_NAMED_IAM capability.

Answer: A

Explanation:

Edit the Trust Policy:

Go to the IAM console in the parent account and locate the role that the EC2 instance needs to assume.

Edit the trust policy of the role to ensure that it correctly allows the sts action for the role ARN in the child account.

Update the Role ARN:

Verify that the target role ARN specified in the trust policy matches the role ARN created by the CloudFormation stack in the child account.

If necessary, update the ARN to reflect the correct role in the child account.

Save and Test:

Save the updated trust policy and ensure there are no syntax errors.

Test the setup by attempting to assume the role from the EC2 instance in the child account. Verify that the instance can successfully assume the role and perform the required actions.

This ensures that the EC2 instance in the child account can assume the role in the parent account, resolving the permission issue.

Reference

AWS IAM Documentation on Trust Policies[51] .

Question: 445

A company is planning to migrate an application from on premises to the AWS Cloud The company will begin the migration by moving the application underlying data storage to AWS The application data is stored on a shared tile system on premises and the application servers connect to the shared file system through SMB A solutions architect must implement a solution that uses an Amazon S3 bucket for shared storage. Until the application is fully migrated and code is rewritten to use native Amazon S3 APIs the application must continue to have access to the data through SMB The solutions architect must migrate the application data to AWS (o its new location while still allowing the on-premises application to access the data

Which solution will meet these requirements?

- A. Create a new Amazon FSx for Windows File Server file system Configure AWS DataSync with one location for the on-premises file share and one location for the new Amazon FSx file system Create a new DataSync task to copy the data from the on-premises file share location to the Amazon FSx file system
- B. Create an S3 bucket for the application Copy the data from the on-premises storage to the S3 bucket
- C. Deploy an AWS Server Migration Service (AWS SMS) VM to the on-premises environment Use AWS SMS to migrate the file storage server from on premises to an Amazon EC2 instance
- D. Create an S3 bucket for the application Deploy a new AWS Storage Gateway file gateway on an on-premises VM Create a new file share that stores data in the S3 bucket and is associated with the file gateway Copy the data from the on-premises storage to the new file gateway endpoint

Answer: D

Explanation:

Create an S3 Bucket:

Log in to the AWS Management Console and navigate to Amazon S3.

Create a new S3 bucket that will serve as the destination for the application data.

Deploy AWS Storage Gateway:

Download and deploy the AWS Storage Gateway virtual machine (VM) on your on-premises environment. This

VM can be deployed on VMware ESXi, Microsoft Hyper-V, or Linux KVM. Configure the File Gateway:

Configure the deployed Storage Gateway as a file gateway. This will enable it to present Amazon S3 buckets as SMB file shares to your on-premises applications.

Create a New File Share:

Within the Storage Gateway configuration, create a new file share that is associated with the S3 bucket you created earlier. This file share will use the SMB protocol, allowing your on-premises applications to access the S3 bucket as if it were a local SMB file share.

Copy Data to the File Gateway:

Use your preferred method (such as robocopy, rsync, or similar tools) to copy data from the on-premises storage to the newly created file gateway endpoint. This data will be stored in the S3 bucket, maintaining accessibility through SMB.

Ensure Secure and Efficient Data Transfer:

AWS Storage Gateway ensures that all data in transit is encrypted using TLS, providing secure data transfer to AWS. It also provides local caching for frequently accessed data, improving access performance for on-premises applications.

This approach allows your existing on-premises applications to continue accessing data via SMB while leveraging the scalability and durability of Amazon S3.

Reference

AWS Storage Gateway Overview [67] .

AWS DataSync and Storage Gateway Hybrid Architecture [66] .

AWS S3 File Gateway Details [68] .

Question: 446

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

Which solutions will meet these requirements'?

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

Answer: C

Explanation:

Create an S3 Bucket:

Navigate to Amazon S3 in the AWS Management Console and create a new S3 bucket to store the game files. Enable static website hosting on this bucket.

Upload Game Files:

Upload the 5 GB game release package to the S3 bucket. Ensure that the files are publicly accessible if

required for download.

Configure Amazon Route 53:

Set up a new domain or subdomain in Amazon Route 53 and point it to the S3 bucket. This allows users to access the game files using a custom URL.

Use Amazon CloudFront:

Create a CloudFront distribution with the S3 bucket as the origin. CloudFront is a content delivery network (CDN) that caches content at edge locations worldwide, improving download performance and reducing latency for users regardless of their location.

Publish the Download URL:

Use the CloudFront distribution URL as the download link for users to access the game files. CloudFront will handle the efficient distribution and caching of the content.

This solution leverages the scalability of Amazon S3 and the performance benefits of CloudFront to provide an optimal download experience for users globally while minimizing costs.

Reference

Amazon CloudFront Documentation

Amazon S3 Static Website Hosting

Question: 447

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

What should a solutions architect do to resolve this issue?

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

Answer: A

Explanation:

Identify Proxy EC2 Instances:

Determine which EC2 instances in the private subnets are running the proxy server software.

Disable Source/Destination Checks:

For each of these EC2 instances, go to the AWS Management Console.

Navigate to the EC2 dashboard, select the instance, and choose "Actions" > "Networking" > "Change Source/Dest.Check".

Disable the source/destination check for these instances.

Disabling source/destination checks allows the EC2 instances to route traffic appropriately, enabling them to function as network appliances or proxies. This ensures that traffic from other instances in the private subnets can be routed through the proxy instances to the internet, meeting the company's security requirements.

Reference

Amazon EC2 User Guide on Source/Destination Checks

Question: 448

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

Which solution will meet these requirements'?

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

Answer: D

Explanation:

Enable AWS Security Hub:

Navigate to the AWS Security Hub console in your management account and enable Security Hub. This process integrates Security Hub with AWS Control Tower, allowing you to manage and monitor security findings across all accounts within your organization.

Designate a Delegated Administrator:

In AWS Organizations, designate one of the AWS accounts as the delegated administrator for Security Hub. This account will have the responsibility to manage and oversee the security posture of all accounts within the organization.

Deploy Controls Across Accounts:

Use AWS Security Hub to automatically enable security controls across all AWS accounts in the organization.

This provides a centralized view of the security state of all accounts and ensures continuous monitoring and compliance.

Utilize AWS Security Hub Features:

Leverage the capabilities of Security Hub to aggregate security alerts, run continuous security checks, and generate findings based on the AWS Foundational Security Best Practices. Security Hub integrates with other AWS services like AWS Config, Amazon GuardDuty, and AWS IAM Access Analyzer to enhance security monitoring and remediation.

By integrating AWS Security Hub with AWS Control Tower and using a delegated administrator account, you can achieve a centralized and comprehensive view of your organization's security posture, facilitating effective management and remediation of security issues.

Reference

AWS Security Hub now integrates with AWS Control Tower [77]

AWS Control Tower and Security Hub Integration [76]

AWS Security Hub Features [79]

Question: 449

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's VPC.

Answer: A,C

Explanation:

Create a Transit Gateway:

Step 1: In the AWS Management Console, navigate to the VPC Dashboard.

Step 2: Select "Transit Gateways" and click on "Create Transit Gateway".

Step 3: Configure the transit gateway by providing a name and setting the options for Amazon side ASN and VPN ECMP support as needed.

Step 4: Attach each of the company's VPCs and relevant subnets to the transit gateway. This centralizes the

network management and simplifies the routing configurations, supporting scalable and flexible network architecture.

Set Up AWS PrivateLink:

Step 1: Create a Network Load Balancer (NLB) in the management VPC that points to the compute resource responsible for license validation.

Step 2: Create an AWS PrivateLink endpoint service pointing to this NLB.

Step 3: Allow each customer's VPC to create an interface endpoint to this PrivateLink service. This setup enables secure and private communication between the customer VPCs and the management VPC, ensuring one-way access from each customer's VPC to the management VPC for license validation.

This combination leverages the benefits of AWS Transit Gateway for scalable and centralized routing, and AWS PrivateLink for secure and private service access, meeting the requirement with minimal operational overhead.

Reference

Amazon VPC-to-Amazon VPC Connectivity Options

AWS PrivateLink - Building a Scalable and Secure Multi-VPC AWS Network Infrastructure Connecting Your VPC to Other VPCs and Networks Using a Transit Gateway

Question: 450

A company needs to use an AWS Transfer Family SFTP-enabled server with an Amazon S3 bucket to receive updates from a third-party data supplier. The data is encrypted with Pretty Good Privacy (PGP) encryption. The company needs a solution that will automatically decrypt the data after the company receives the data. A solutions architect will use a Transfer Family managed workflow. The company has created an IAM service role by using an IAM policy that allows access to AWS Secrets Manager and the S3 bucket. The role's trust relationship allows the transfer.amazonaws.com service to assume the role. What should the solutions architect do next to complete the solution for automatic decryption?

- A. Store the PGP public key in Secrets Manager. Add a nominal step in the Transfer Family managed workflow to decrypt files. Configure PGP encryption parameters in the nominal step. Associate the workflow with the Transfer Family server.
- B. Store the PGP private key in Secrets Manager. Add an exception-handling step in the Transfer Family managed workflow to decrypt files. Configure PGP encryption parameters in the exception handler. Associate the workflow with the SFTP user.
- C. Store the PGP private key in Secrets Manager. Add a nominal step in the Transfer Family managed workflow to decrypt files. Configure PGP decryption parameters in the nominal step. Associate the workflow with the Transfer Family server.
- D. Store the PGP public key in Secrets Manager. Add an exception-handling step in the Transfer Family managed workflow to decrypt files. Configure PGP decryption parameters in the exception handler. Associate the workflow with the SFTP user.

Answer: C

Explanation:

Store the PGP Private Key:

Step 1: In the AWS Management Console, navigate to AWS Secrets Manager.

Step 2: Store the PGP private key in Secrets Manager. Ensure the key is encrypted and properly secured.

Set Up the Transfer Family Managed Workflow:

Step 1: In the AWS Transfer Family console, create a new managed workflow.

Step 2: Add a nominal step to the workflow that includes the decryption of the files. Configure this step with the PGP decryption parameters, referencing the PGP private key stored in Secrets Manager. Step 3: Associate this workflow with the Transfer Family SFTP server, ensuring that incoming files are automatically decrypted upon receipt.

This solution ensures that the data is securely decrypted as it is transferred from the SFTP server to the S3 bucket, automating the decryption process and leveraging AWS Secrets Manager for key management.

Reference

AWS Transfer Family Documentation

Using AWS Secrets Manager for Managing Secrets

AWS Transfer Family Managed Workflows

Question: 451

A delivery company is running a serverless solution in the AWS Cloud. The solution manages user data, delivery information, and past purchase details. The solution consists of several microservices. The central user service stores sensitive data in an Amazon DynamoDB table. Several of the other microservices store a copy of parts of the sensitive data in different storage services. The company needs the ability to delete user information upon request. As soon as the central user service deletes a user, every other microservice must also delete its copy of the data immediately. Which solution will meet these requirements?

- A. Activate DynamoDB Streams on the DynamoDB table. Create an AWS Lambda trigger for the DynamoDB stream that will post events about user deletion in an Amazon Simple Queue Service (Amazon SQS) queue. Configure each microservice to poll the queue and delete the user from the DynamoDB table.
- B. Set up DynamoDB event notifications on the DynamoDB table. Create an Amazon Simple Notification Service (Amazon SNS) topic as a target for the DynamoDB event notification. Configure each microservice to subscribe to the SNS topic and to delete the user from the DynamoDB table.
- C. Configure the central user service to post an event on a custom Amazon EventBridge event bus when the company deletes a user. Create an EventBridge rule for each microservice to match the user deletion event pattern and invoke logic in the microservice to delete the user from the DynamoDB table.
- D. Configure the central user service to post a message on an Amazon Simple Queue Service (Amazon SQS) queue when the company deletes a user. Configure each microservice to create an event filter on the SQS queue and to delete the user from the DynamoDB table.

Answer: C

Explanation:

Set Up EventBridge Event Bus:

Step 1: Open the Amazon EventBridge console and create a custom event bus. This bus will be used to handle user deletion events.

Step 2: Name the event bus appropriately (e.g., user-deletion-bus).

Post Events on User Deletion:

Step 1: Modify the central user service to post an event to the custom EventBridge event bus whenever a user is deleted.

Step 2: Ensure the event includes relevant details such as the user ID and any other necessary metadata.

Create EventBridge Rules for Microservices:

Step 1: For each microservice that needs to delete user data, create a new rule in EventBridge that triggers on the user deletion event.

Step 2: Define the event pattern to match the user deletion event. This pattern should include the event details posted by the central user service.

Invoke Microservice Logic:

Step 1: Configure the EventBridge rule to invoke a target, such as an AWS Lambda function, which contains the logic to delete the user data from the microservice's data store.

Step 2: Each microservice should have its Lambda function or equivalent logic to handle the deletion of user data upon receiving the event.

Using Amazon EventBridge ensures a scalable, reliable, and decoupled approach to handle the deletion of user data across multiple microservices. This setup allows each microservice to independently process user deletion events without direct dependencies on other services. **Reference**

[AWS EventBridge Documentation](#)

[DynamoDB Streams and AWS Lambda Triggers](#)

[Implementing the Transactional Outbox Pattern with EventBridge Pipes\(AWS Documentation\)\(Amazon Web Services, Inc.\)\(Amazon Web Services, Inc.\)\(AWS Documentation\)\(AWS Cloud Community\).](#)

Question: 452

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 the company WAN to 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 AWS Direct Connect Connections:

Step 1: Set up two AWS Direct Connect (DX) connections from the company headquarters to a chosen AWS Region. This provides a redundant and high-availability setup to ensure continuous connectivity. Step 2: Ensure that these DX connections terminate in a specific Direct Connect location associated with the chosen AWS Region.

Use Company WAN:

Step 1: Configure the company's global WAN to route traffic through the established Direct Connect connections.

Step 2: This setup ensures that all traffic between the company's headquarters and AWS does not traverse the public internet, maintaining compliance with security requirements.

Set Up Direct Connect Gateway:

Step 1: Create a Direct Connect Gateway in the AWS Management Console. This gateway allows you to connect your Direct Connect connections to multiple VPCs across different AWS Regions.

Step 2: Associate the Direct Connect Gateway with the VPCs in the various Regions where your critical data is stored. This enables access to data in multiple Regions through a single Direct Connect connection.

By using Direct Connect and Direct Connect Gateway, the company can achieve secure, reliable, and cost-effective access to data stored across multiple AWS Regions without using the public internet, ensuring compliance with industry regulations.

Reference

AWS Direct Connect Documentation

Building a Scalable and Secure Multi-VPC AWS Network Infrastructure(AWS Documentation)(AWS Documentation).

Question: 453

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 CloudFormation 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 CloudFormation templates to provision Amazon EC2 instances and use AWS CloudFormation Init commands to mount the Amazon FSx file shares

Answer: B

Explanation:

Set Up AWS Elastic Disaster Recovery:

Navigate to the AWS Elastic Disaster Recovery (DRS) console.

Configure the Elastic Disaster Recovery service to replicate your on-premises VMware vSphere VMs to Amazon EC2 instances. This involves installing the AWS Replication Agent on your VMs.

Configure Replication Settings:

Define the replication settings, including the Amazon EC2 instance type and the Amazon EBS volume configuration. Ensure that the replication frequency meets your Recovery Point Objective (RPO) of 5 minutes.

Monitor Data Replication:

Monitor the initial data replication process in the Elastic Disaster Recovery console. Once the initial sync is complete, the status should show as "Healthy" indicating that the data replication is up-to-date and within the RPO requirements.

Disaster Recovery (Failover):

In the event of a disaster, initiate a failover from the Elastic Disaster Recovery console. This will launch the replicated Amazon EC2 instances using the Amazon EBS volumes with the latest data. **Failback Process:**

Once the on-premises environment is restored, perform a failback operation to synchronize the data from AWS back to your on-premises VMware environment. Use the failback client provided by AWS Elastic Disaster Recovery to ensure data consistency and minimal downtime during the failback process.

Using AWS Elastic Disaster Recovery provides a low-overhead, automated solution for disaster recovery that ensures minimal data loss and meets the RPO requirement of 5 minutes (Amazon Web Services, Inc.) (AWS Documentation).

Question: 454

A company provides a centralized Amazon EC2 application hosted in a single shared VPC. The centralized application must be accessible from client applications running in the VPCs of other business units. The centralized application front end is configured with a Network Load Balancer (NLB) for scalability. Up to 10 business unit VPCs will need to be connected to the shared VPC. Some of the business unit VPC CIDR blocks overlap with the shared VPC and some overlap with each other. Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only.

Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

- A. Create an AWS Transit Gateway. Attach the shared VPC and the authorized business unit VPCs to the transit gateway. Create a single transit gateway route table and associate it with all of the attached VPCs. Allow automatic propagation of routes from the attachments into the route table. Configure VPC routing tables to send traffic to the transit gateway.
- B. Create a VPC endpoint service using the centralized application NLB and enable the option to require endpoint acceptance. Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service. Accept authorized endpoint requests from the endpoint service console.
- C. Create a VPC peering connection from each business unit VPC to the shared VPC. Accept the VPC peering

connections from the shared VPC console Configure VPC routing tables to send traffic to the VPC peering connection

D. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPCs Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC Configure VPC routing tables to send traffic to the VPN connection

Answer: B

Explanation:

Create VPC Endpoint Service:

In the shared VPC, create a VPC endpoint service using the Network Load Balancer (NLB) that fronts the centralized application.

Enable the option to require endpoint acceptance to control which business unit VPCs can connect to the service.

Set Up VPC Endpoints in Business Unit VPCs:

In each business unit VPC, create a VPC endpoint that points to the VPC endpoint service created in the shared VPC.

Use the service name of the endpoint service created in the shared VPC for configuration. **Accept Endpoint Requests:**

From the VPC endpoint service console in the shared VPC, review and accept endpoint connection requests from authorized business unit VPCs. This ensures that only authorized VPCs can access the centralized application.

Configure Routing:

Update the route tables in each business unit VPC to direct traffic destined for the centralized application through the VPC endpoint.

This solution ensures secure, private connectivity between the business unit VPCs and the shared VPC, even if there are overlapping CIDR blocks. It leverages AWS PrivateLink and VPC endpoints to provide scalable and controlled access(AWS Documentation)(Amazon Web Services, Inc.).

Question: 455

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 logs and block any requests that come from unauthenticated users.

Answer: A

Explanation:

Integrate ALB with OIDC IdP:

In the AWS Management Console, navigate to the Application Load Balancer (ALB) settings. Configure the ALB to use the OpenID Connect (OIDC) IdP for authentication. This ensures that all requests routed through the ALB are authenticated using the IdP.

Set Up Authentication Rules:

Create a listener rule on the ALB that requires authentication. This rule will forward requests to the IdP for user authentication before allowing access to the backend services.

Restrict Unauthenticated Access:

Ensure the ALB only forwards requests to backend services if the user is authenticated. Unauthenticated requests should be blocked or redirected to the IdP for authentication.

Update CloudFront Configuration:

Modify the CloudFront distribution to forward authenticated requests to the ALB. Ensure that the ALB and API Gateway accept only requests coming through the CloudFront distribution to enforce consistent authentication and security.

By enforcing authentication at the ALB level, you ensure that all backend services are accessed only by authenticated users, enhancing the overall security of the web application

Question: 456

A medical company is running a REST API on a set of Amazon EC2 instances The EC2 instances run in an Auto Scaling group behind an Application Load Balancer (ALB) The ALB runs in three public subnets, and the EC2 instances run in three private subnets The company has deployed an Amazon CloudFront distribution that has the ALB as the only origin

Which solution should a solutions architect recommend to enhance the origin security?

- A. Store a random string in AWS Secrets Manager Create an AWS Lambda function for automatic secret rotation Configure CloudFront to inject the random string as a custom HTTP header for the origin request Create an AWS WAF web ACL rule with a string match rule for the custom header Associate the web ACL with the ALB
- B. Create an AWS WAF web ACL rule with an IP match condition of the CloudFront service IP address ranges Associate the web ACL with the ALB Move the ALB into the three private subnets
- C. Store a random string in AWS Systems Manager Parameter Store Configure Parameter Store automatic rotation for the string Configure CloudFront to inject the random string as a custom HTTP header for the origin request Inspect the value of the custom HTTP header, and block access in the ALB
- D. Configure AWS Shield Advanced. Create a security group policy to allow connections from CloudFront service IP address ranges. Add the policy to AWS Shield Advanced, and attach the policy to the ALB

Answer: A

Explanation:

Store Secret in AWS Secrets Manager:

Create a random string in AWS Secrets Manager to be used as a custom HTTP header value.

Set Up Automatic Rotation:

Implement a Lambda function to handle automatic rotation of the secret in AWS Secrets Manager, ensuring the header value remains secure.

Configure CloudFront Custom Header:

In the CloudFront distribution settings, configure an origin custom header with the name and value from AWS Secrets Manager. This header will be included in requests forwarded to the ALB. Create AWS WAF Web ACL:

Create a Web ACL in AWS WAF with a string match rule to allow requests that include the custom header with the correct value.

Associate the Web ACL with the ALB to filter incoming traffic based on the custom header.

By using this method, you can ensure that only requests coming through CloudFront (which injects the custom header) can reach the ALB, enhancing the origin security

Question: 457

A company has multiple lines of business (LOBs) that toll up to the parent company. The company has asked its solutions architect to develop a solution with the following requirements

- Produce a single AWS invoice for all of the AWS accounts used by its LOBs.
- The costs for each LOB account should be broken out on the invoice
- Provide the ability to restrict services and features in the LOB accounts, as defined by the company's governance policy
- Each LOB account should be delegated full administrator permissions regardless of the governance policy

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

- A. Use AWS Organizations to create an organization in the parent account for each LOB Then invite each LOB account to the appropriate organization
- B. Use AWS Organizations to create a single organization in the parent account Then, invite each LOB's AWS account to join the organization.
- C. Implement service quotas to define the services and features that are permitted and apply the quotas to each LOB. as appropriate
- D. Create an SCP that allows only approved services and features then apply the policy to the LOB accounts
- E. Enable consolidated billing in the parent account's billing console and link the LOB accounts

Answer: B,E

Explanation:

Create AWS Organization:

In the AWS Management Console, navigate to AWS Organizations and create a new organization in the parent account.

Invite LOB Accounts:

Invite each Line of Business (LOB) account to join the organization. This allows centralized management and governance of all accounts.

Enable Consolidated Billing:

Enable consolidated billing in the billing console of the parent account. Link all LOB accounts to ensure a single consolidated invoice that breaks down costs per account.

Apply Service Control Policies (SCPs):

Implement Service Control Policies (SCPs) to define the services and features permitted for each LOB account as per the governance policy, while still delegating full administrative permissions to the LOB accounts.

By consolidating billing and using AWS Organizations, the company can achieve centralized billing and governance while maintaining independent administrative control for each LOB account

Question: 458

A company is migrating its blog platform to AWS. The company's on-premises servers connect to AWS through an AWS Site-to-Site VPN connection. The blog content is updated several times a day by multiple authors and is served from a file share on a network-attached storage (NAS) server. The company needs to migrate the blog platform without delaying the content updates. The company has deployed Amazon EC2 instances across multiple Availability Zones to run the blog platform behind an Application Load Balancer. The company also needs to move 200 TB of archival data from its on-premises servers to Amazon S3 as soon as possible.

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

- A. Create a weekly cron job in Amazon EventBridge. Use the cron job to invoke an AWS Lambda function to update the EC2 instances from the NAS server.
- B. Configure an Amazon Elastic Block Store (Amazon EBS) Multi-Attach volume for the EC2 instances to share for content access. Write code to synchronize the EBS volume with the NAS server weekly.
- C. Mount an Amazon Elastic File System (Amazon EFS) file system to the on-premises servers to act as the NAS server. Copy the blog data to the EFS file system. Mount the EFS file system to the EC2 instances to serve the content.
- D. Order an AWS Snowball Edge Storage Optimized device. Copy the static data artifacts to the device. Ship the device to AWS.
- E. Order an AWS Snowcone SSD device. Copy the static data artifacts to the device. Ship the device to AWS.

Answer: C,D

Explanation:

Comprehensive and Detailed in Depth

C is correct because Amazon EFS offers a scalable, elastic file system that can be mounted on both on-premises servers (via AWS Direct Connect or VPN) and EC2 instances. This allows real-time access to shared content with no migration downtime.

D is correct because AWS Snowball Edge Storage Optimized devices are designed for transferring petabyte-scale data (up to 80 TB per device). It's a best practice for moving 200 TB of archival data when speed and bandwidth are constraints.

Reference:

Amazon EFS - Mounting on-premises

AWS Snowball Edge - Overview

Question: 459

A company is designing its network configuration in the AWS Cloud. The company uses AWS Organizations to manage a multi-account setup. The company has three OUs. Each OU contains more than 100 AWS accounts. Each account has a single VPC, and all the VPCs in each OU are in the same AWS Region.

The CIDR ranges for all the AWS accounts do not overlap. The company needs to implement a solution in which VPCs in the same OU can communicate with each other but cannot communicate with VPCs in other OUs.

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

- A. Create an AWS CloudFormation stack set that establishes VPC peering between accounts in each OU. Provision the stack set in each OU.
- B. In each OU, create a dedicated networking account that has a single VPC. Share this VPC with all the other accounts in the OU by using AWS Resource Access Manager (AWS RAM). Create a VPC peering connection between the networking account and each account in the OU.
- C. Provision a transit gateway in an account in each OU. Share the transit gateway across the organization by using AWS Resource Access Manager (AWS RAM). Create transit gateway VPC attachments for each VPC.
- D. In each OU, create a dedicated networking account that has a single VPC. Establish a VPN connection between the networking account and the other accounts in the OU. Use third-party routing software to route transitive traffic between the VPCs.

Answer: C

Explanation:

Comprehensive and Detailed in Depth C is correct because AWS Transit Gateway is the most scalable and efficient way to interconnect hundreds of VPCs. By deploying one transit gateway per OU and sharing it with AWS RAM, each OU can isolate its network traffic and maintain internal communication without affecting or exposing other OUs.

Reference:

AWS Transit Gateway Best Practices Using RAM with Transit Gateway

Question: 460

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 use an Amazon EC2 instance with IDE access to modify and test applications. The developers want to implement CI/CD with AWS CodePipeline with the following requirements:

Use AWS CodeCommit for source control.

Automate unit testing and security scanning.

Alert developers when unit tests fail.

Toggle application features and allow lead developer approval before deployment.

Which solution will meet these requirements?

- A. Use AWS CodeBuild for testing and scanning. Use EventBridge and SNS for alerts. Use AWS CDK with a manifest to toggle features. Use a manual approval stage.
- B. Use Lambda for testing and alerts. Use AWS Amplify plugins for feature toggles. Use SES for manual approval.
- C. Use Jenkins and SES for alerts. Use nested CloudFormation stacks for features. Use Lambda for approvals.
- D. Use CodeDeploy for testing and scanning. Use CloudWatch alarms and SNS. Use Docker images for features and AWS CLI for toggles.

Answer: A

Explanation:

Comprehensive and Detailed in Depth

A is correct because AWS CodeBuild is a fully managed build service that supports running tests and security scans. EventBridge can detect test failures and send notifications via SNS. AWS CDK with manifest files is a robust way to manage feature toggles. CodePipeline supports manual approval stages to allow lead developer intervention.

Reference:

[AWS CodePipeline Approvals](#)

[AWS CDK with feature flags](#)

Question: 461

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:

Comprehensive and Detailed in Depth

A is correct because write throttling on DynamoDB means WCU (write capacity units) are insufficient. Increasing them will reduce the error rate.

D is correct because introducing Kinesis allows for efficient, high-throughput ingestion and batch processing, reducing the number of Lambda invocations and overall load.

Reference:

DynamoDB Throttling
Using Kinesis with Lambda

Question: 462

A company runs a serverless ecommerce application on AWS. The application uses API Gateway to invoke Java Lambda functions that connect to an Amazon RDS for MySQL database. During a sale event, traffic spikes caused slow performance and DB connection failures.

Which solution will improve performance with the LEAST application change?

- A. Move DB connection outside Lambda handler and increase provisioned concurrency.
- B. Use RDS Proxy. Store DB credentials in Secrets Manager. Update Lambda to use RDS Proxy. Increase provisioned concurrency.
- C. Increase max_connections parameter in a custom DB parameter group and reboot. Increase reserved concurrency.
- D. Use RDS Proxy and Secrets Manager. Increase reserved concurrency.

Answer: B

Explanation:

Comprehensive and Detailed in Depth

B is correct because RDS Proxy allows for connection pooling and improved management of DB connections. Lambda + RDS often runs into connection exhaustion during high concurrency unless RDS Proxy is used. Provisioned concurrency prevents cold starts.

Reference:

RDS Proxy Overview
Provisioned Concurrency

Question: 463

A company has an application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The application is in an AWS account that has AWS CloudTrail enabled. The company restricts access to the application by adding the IP addresses of end users to a security group that is associated with the ALB.

The company is developing an AWS Lambda function to determine if the allowed IP addresses have accessed the application recently. If an allowed IP address has not accessed the application in the last 90 days, the Lambda function will remove the IP address from the security group.

The company needs to implement the functionality for the Lambda function to check the IP addresses.

Which combination of steps will provide this functionality MOST cost-effectively? (Select TWO.)

- A. For the VPC that contains the ALB, configure VPC flow logs to be sent to a log group in Amazon CloudWatch Logs.
- B. Enable access logging on the ALB. Create an Amazon Athena table to query the ALB access logs.
- C. Program the Lambda function to check when each allowed IP address from the security group last appeared in the VPC flow logs.
- D. Program the Lambda function to check when each allowed IP address from the security group last appeared in the ALB access logs.
- E. Program the Lambda function to check when each allowed IP address from the security group last appeared in the CloudTrail logs.

Answer: B,D

Explanation:

B is correct because ALB access logs contain the exact IP addresses of all incoming client requests, and storing these logs in S3 enables historical analysis. Using Amazon Athena to query them offers an efficient, serverless way to check when each IP last accessed the application.

D is correct because parsing ALB access logs gives direct evidence of activity from IP addresses — ideal for implementing the "last accessed in 90 days" rule.

A and C refer to VPC flow logs, which only capture network metadata, not HTTP request-level details like the ALB logs do.

E is incorrect — CloudTrail logs AWS API actions, not web requests through an ALB.

Reference:

Access logs for your Application Load Balancer

Querying ALB logs with Amazon Athena

Question: 464

A company has separate AWS accounts for each of its departments. The accounts are in OUs that are in an organization in AWS Organizations. The IT department manages a private certificate authority (CA) by using AWS Private Certificate Authority in its account.

The company needs a solution to allow developer teams in the other departmental accounts to access the private CA to issue certificates for their applications. The solution must maintain appropriate security boundaries between accounts.

Which solution will meet these requirements?

- A. Create an AWS Lambda function in the IT account. Program the Lambda function to use the AWS Private CA API to export and import a private CA certificate to each department account. Use Amazon EventBridge to invoke the Lambda function on a schedule.
- B. Create an IAM identity-based policy that allows cross-account access to AWS Private CA. In the IT account, attach this policy to the private CA. Grant access to AWS Private CA by using the AWS Private CA API.
- C. In the organization's management account, create an AWS CloudFormation stack to set up a resource-based delegation policy.
- D. Use AWS Resource Access Manager (AWS RAM) in the IT account to enable sharing in the organization. Create a resource share. Add the private CA resource to the resource share. Grant the department OUs access to the shared CA.

Answer: D

Explanation:

D is correct because AWS Private CA supports resource sharing through AWS RAM, which allows you to share the CA across accounts in your AWS Organization securely. It ensures the CA private key remains secure and is never exported.

A is invalid because you cannot export a CA's private key, and importing/exporting CAs this way is unsupported.

B is incorrect — IAM policies alone are not sufficient to share Private CAs across accounts.

C is not applicable because resource-based delegation policies do not apply to AWS Private CA. **Reference:**

Share your private CA using AWS RAM

Using AWS RAM for cross-account sharing

Question: 465

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.

Answer: C

Explanation:

C is correct because:

AWS WAF integrates natively with API Gateway and protects against common web exploits (e.g., SQL injection, XSS).

Amazon Inspector can scan the legacy EC2 instance for known vulnerabilities.

Amazon GuardDuty is a continuous security monitoring service that detects threats but does not block traffic (B and D are incorrect because GuardDuty doesn't block).

Reference:

[AWS WAF Overview](#)

[Amazon Inspector Overview](#)

[Amazon GuardDuty Overview](#)

Question: 466

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

Must support iSCSI access from on-premises servers.

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

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

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

Answer: C

Explanation:

C is correct because the Storage Gateway Volume Gateway in cache mode supports iSCSI block storage, which is exactly what legacy applications need. It also caches frequently accessed data locally for low-latency performance and integrates with AWS Backup for point-in-time copies. A is incorrect — tape gateway is designed for archival backup only.

B and D are incorrect because FSx and file gateway offer file (NFS/SMB) storage, not block/iSCSI support.

Reference:

[Volume Gateway Overview](#)

[AWS Backup Integration](#)

Question: 467

A company completed a successful Amazon Workspaces proof of concept. They now want to make Workspaces highly available across two AWS Regions. Workspaces are deployed in the failover Region. A hosted zone is available in Amazon Route 53.

What should the solutions architect do?

- A. Create a connection alias in the primary Region and in the failover Region. Associate each with a directory in its Region. Create a Route 53 failover routing policy with Evaluate Target Health = Yes.
- B. Create a connection alias in both Regions. Associate both with a directory in the primary Region. Use a Route 53 multivalued answer routing policy.
- C. Create a connection alias in the primary Region. Associate with the directory in the primary Region. Use Route 53 weighted routing.
- D. Create a connection alias in the primary Region. Associate it with the directory in the failover Region. Use Route 53 failover routing with Evaluate Target Health = Yes.

Answer: A

Explanation:

A is correct because AWS recommends using one connection alias per Region, associated with each directory. Then, configure a Route 53 failover policy so that if the primary Region becomes unhealthy, users are directed to the failover Region automatically. "Evaluate Target Health" ensures automatic detection and failover.

Reference:

Amazon Workspaces Cross-Region Resilience

Route 53 Failover Routing

Question: 468

A company hosts an application that uses several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB). During the initial startup of the EC2 instances, the EC2 instances run user data scripts to download critical content for the application from an Amazon S3 bucket.

The EC2 instances are launching correctly. However, after a period of time, the EC2 instances are terminated with the following error message:

"An instance was taken out of service in response to an ELB system health check failure." The only recent change to the deployment is that the company added a large amount of critical content to the S3 bucket.

What should a solutions architect do so that the production environment can deploy successfully?

- A. Increase the size of the EC2 instances.
- B. Increase the health check timeout for the ALB.
- C. Change the health check path for the ALB.
- D. Increase the health check grace period for the Auto Scaling group.

Answer: D

Explanation:

D is correct because the health check grace period defines how long Auto Scaling waits after launching an instance before checking its health status. With the larger content added to the S3 bucket, the instance initialization (via user data) is taking longer. Increasing the grace period allows the instance time to complete startup tasks before it's marked as unhealthy.

Option A would not necessarily reduce startup time — the issue is likely network latency or the size of the content.

Option B only affects the timeout for health check responses, not the delay before they begin. Option C is not applicable unless the health check path itself is invalid (which isn't mentioned). **Reference:**

Auto Scaling Health Checks

Question: 469

A solutions architect is importing a VM from an on-premises environment by using the Amazon EC2 VM Import feature. The imported instance has a public IP and runs in a public subnet in a VPC. However, the instance does not appear in the AWS Systems Manager (SSM) console as a managed instance. Which combination of steps should the architect take to resolve the issue? (Select TWO.)

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

Answer: A,B

Explanation:

A is correct because SSM Agent must be installed and actively running on the instance for it to be managed by Systems Manager. Imported VMs typically do not have the agent installed by default. B is also correct because an IAM instance profile with at least the AmazonSSMManagedInstanceCore policy must be attached to the instance to allow communication with SSM.

Option C is unnecessary if the instance has internet access (via public IP). Option D relates to migration assessments, not Systems Manager.

Option E is good practice but not required for the instance to appear as managed. Reference:

Managed Instance Requirements
SSM Agent Installation

Question: 470

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 use default AWS-managed KMS encryption keys. The backups must be stored in a WORM (Write Once Read Many) backup account. What is the correct approach to enable cross-account backup?

- A. Restore the databases with customer-managed KMS keys and use AWS Backup with cross-account vault sharing.
- 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:

A is correct because default AWS-managed KMS keys cannot be shared across accounts. Therefore, the only way to enable cross-account backup is to restore each RDS instance using a customer-managed key (CMK) and share that key with the central account.

Then, AWS Backup can perform cross-account backup operations using AWS Organizations trusted access and backup policies.

Option B is incorrect due to the inability to share default keys.

Option C and D are technically invalid — RDS snapshots cannot be decrypted or re-encrypted manually.

Reference:

Cross-Account Backups with AWS Backup

RDS Snapshot Sharing and KMS

Question: 471

A company mandates that all internal AWS communications use private IPs. A solutions architect created interface VPC endpoints for public AWS services like S3. However, service names are still resolving to public IP addresses, and the internal apps cannot connect.

What should the architect do to resolve this issue?

- A. Update the subnet route table with a route to the interface endpoint.
- B. Enable the private DNS option on the VPC attributes.
- C. Configure the security group on the interface endpoint to allow access.
- D. Configure a private hosted zone with conditional forwarding.

Answer: B

Explanation:

B is correct. When Private DNS is enabled for an interface endpoint, AWS automatically updates the public service DNS names (e.g., s3.amazonaws.com) to resolve to private IPs inside the VPC. This ensures all traffic stays within the AWS network and complies with the private IP policy.

A is not necessary — VPC endpoints already route via local network.

C only affects security group rules, not DNS resolution.

D is unnecessary unless using custom DNS resolution systems.

Reference:

VPC Endpoints and Private DNS

Question: 472

A company is modernizing a legacy .NET Framework application backed by SQL Server. Requirements: Containerize into microservices.

Control OS patches and storage.

Add load balancing.

Ensure high availability. Which solution meets all of these with minimal refactoring?

- A. Use App2Container to deploy on ECS EC2 with ALB and RDS for SQL Server.
- B. Use App2Container on ECS EC2 with NLB and Aurora MySQL.
- C. Use Porting Assistant and EKS with Fargate and Aurora MySQL.
- D. Use Porting Assistant and EKS with Fargate and RDS SQL Server.

Answer: A

Explanation:

A is correct because:

App2Containers supports packaging .NET Framework apps into containers without porting to .NET

Core.

ECS with EC2 gives full control over the OS and patching.

ALB handles microservice-level load balancing.

RDS for SQL Server with Multi-AZ ensures high availability.

Options B, C, and D involve Aurora MySQL (incompatible with SQL Server features) or require .NET Core, which involves more significant application changes.

Reference:

App2Container Overview

ECS EC2 vs Fargate

Question: 473

A company has an application that uses AWS Key Management Service (AWS KMS) to encrypt and decrypt data. The application stores data in an Amazon S3 bucket in an AWS Region. Company security policies require that the data is encrypted before being uploaded to S3, and decrypted when read. The S3 bucket is replicated to other AWS Regions.

A solutions architect must design a solution so that the application can encrypt and decrypt data across Regions using the same key.

Options:

- A. Create a KMS multi-Region primary key. Use it to create KMS multi-Region replica keys in each Region. Update application code to use the replica key in each Region.
- B. Create a new customer-managed KMS key in each additional Region. Update application code to use the key in each Region.
- C. Use AWS Private CA to issue TLS certificates and replicate them with AWS RAM.
- D. Export the KMS key material to Systems Manager Parameter Store in each Region. Update the app to use those.

Answer: A

Explanation:

AWSKMS multi-Region keys allow encryption in one Region and decryption in another. This is specifically built for cross-Region replication scenarios, like Amazon S3 Cross-Region Replication (CRR).

With a multi-Region key, you create a primary key in one Region and replica keys in others. These replicas share the same key material and key ID.

Option B would fail because different keys = different IDs = decryption mismatch. Option C is unrelated — Private CA is for certificates, not symmetric key encryption. Option D is insecure and violates AWS KMS key management best practices.

Reference: <https://docs.aws.amazon.com/kms/latest/developerguide/multi-region-keys-overview.html>

Question: 474

A company runs workloads on EC2 in multiple VPCs in a single Region. They also have an on-premises DNS server (via Direct Connect). All EC2 instances must resolve internal.company.com using private communication.

What should a solutions architect do? (Select THREE.)

Options:

- A. Create an Amazon Route 53 inbound endpoint in all workload VPCs.
- B. Create a Route 53 outbound endpoint in one VPC.
- C. Create a Route 53 forwarding rule to forward internal.company.com to the on-prem DNS.
- D. Create a Route 53 rule with the System type.
- E. Associate the rule with all VPCs.
- F. Associate the rule only with the VPC that has the outbound endpoint.

Answer: B,C,E

Explanation:

B: Outbound Resolver endpoints enable EC2 instances to send DNS queries to external resolvers (i.e., on-prem DNS servers).

C: A forwarding rule must be created to forward the internal.company.com domain to the on-prem resolver.

E: You must associate the rule with every VPC that needs to forward the queries.

Incorrect:

A: Inbound endpoints are for receiving DNS queries from on-prem to AWS, not vice versa.

D: There's no such thing as a "System" rule type.

F: Limiting to one VPC won't meet the requirement.

Reference: <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver.html>

Question: 475

A company uses AWS Organizations and tags every resource with a BusinessUnit tag. They want to allocate

cloud costs by business unit and visualize them.

Options:

- A. Activate BusinessUnit cost allocation tag in the management account. Create a CUR to S3. Use Athena + QuickSight for reporting.
- B. Create cost allocation tags in each member account. Use CloudWatch Dashboards.
- C. Create cost allocation tags in the management account. Deploy CURs per account.
- D. Use tags and CUR per account. Visualize with QuickSight from management account.

Answer: A

Explanation:

A is the correct best-practice approach:

Activate cost allocation tags in management/payer account.

Enable AWS Cost and Usage Report (CUR) to dump usage to S3.

Query CUR with Amazon Athena, and visualize with Amazon QuickSight.

Other options either duplicate CURs (C, D) or misuse tools (B: CloudWatch is not a cost analysis tool).

Reference: <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html> <https://docs.aws.amazon.com/cur/latest/userguide/what-is-cur.html>

Question: 476

A company uses IAM Identity Center for data scientist access. Each user should be able to access only their own data in an S3 bucket. The company also needs to generate monthly access reports per user. Options:

- A. Use IAM Identity Center permission sets to allow S3 access scoped to userName tag.
- B. Use a shared IAM Identity Center role for all users and bucket policy.
- C. Use AWS CloudTrail to log S3 data events, query via Athena.
- D. Use CloudTrail management events to CloudWatch, then use Athena.
- E. Use S3 access logs and S3 Select for reporting.

Answer: A,C

Explanation:

A: Use dynamic IAM policies with {aws:PrincipalTag/userName} to enforce prefix-level access control— i.e., bucket/userA/*, bucket/userB/*.

C: Enable CloudTrail data events to capture object-level access and query them with Athena. This is the AWS-recommended way to audit per-user object access.

Incorrect:

B doesn't provide user isolation.

D only captures management events, not object-level data access.

E is legacy, inefficient, and not structured for per-user auditing.

Reference: https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_variables.html

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/logging-data-events.html>

Question: 477

A SaaS web app runs on EC2 Linux behind an ALB. It stores user sessions in an RDS Multi-AZ database. During high traffic, the app suffers latency due to session read/write.

What is the best way to reduce session latency?

Options:

- A. Store session data in Amazon S3.
- B. Use FSx for Windows and mount it.
- C. Use Multi-Attach EBS volumes.
- D. Use ElastiCache for Redis to store sessions.

Answer: D

Explanation:

Dis the AWS best practice for session storage: Use ElastiCache for Redis — a fast, in-memory data store that handles high throughput with microsecond latency.

It's highly scalable, fault-tolerant, and optimized for temporary, fast-access session data.

Incorrect:

A: S3 is slow and object-based — not for session I/O.

B: FSx is Windows-only and not ideal for this use case.

C: EBS Multi-Attach has limitations, complexity, and is not suitable for high-performance shared memory.

Reference: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/WhatIs.html>

Question: 478

A company is migrating a monolithic on-premises .NET Framework production application to AWS. Application demand will grow exponentially in the next 6 months. The company must ensure that the application can scale appropriately.

The application currently connects to a Microsoft SQL Server transactional database. The company has well-documented source code for the application. Some business logic is contained within stored procedures.

A solutions architect must recommend a solution to redesign the application to meet the growth in demand. Which solution will meet this requirement MOST cost-effectively?

- A. Use Amazon API Gateway APIs and Amazon EC2 Spot Instances to rehost the application with a scalable microservices architecture. Deploy the EC2 instances in a cluster placement group. Configure EC2 Auto Scaling. Store the data and stored procedures in Amazon RDS for SQL Server.
- B. Use AWS Application Migration Service to migrate the application to AWS Elastic Beanstalk. Deploy Elastic Beanstalk packages to configure and deploy the application as microservices. Deploy Elastic Beanstalk across multiple Availability Zones and configure auto scaling. Store the data and stored procedures in Amazon RDS for MySQL.

- C. Migrate the applications by using AWS App2Container. Use AWS Fargate in multiple AWS Regions to host the containers. Use Amazon API Gateway APIs and AWS Lambda functions to call the containers. Store the data and stored procedures in Amazon DynamoDB Accelerator (DAX).
- D. Use Amazon API Gateway APIs and AWS Lambda functions to decouple the application into microservices. Use the AWS Schema Conversion Tool (AWS SCT) to review and modify the stored procedures. Store the data in Amazon Aurora Serverless v2.

Answer: D

Explanation:

D is correct because this solution modernizes the application into a serverless architecture using API Gateway and Lambda for scalable microservices. Aurora Serverless v2 supports SQL workloads and auto-scales based on demand. AWS SCT allows migration of SQL Server stored procedures to Aurora-compatible formats. This setup ensures cost efficiency, scalability, and minimal manual intervention. A rehosts but doesn't refactor into microservices.

B uses MySQL, which might not support the SQL Server-specific stored procedures fully.

C adds unnecessary complexity and loses relational database functionality by using DAX.

Reference: <https://docs.aws.amazon.com/aurora/latest/aurora-serverless/aurora-serverless.html>

Question: 479

A company is running a large containerized workload in the AWS Cloud using Amazon ECS. The development team recently started using AWS Fargate instead of EC2 in the ECS cluster. The company is worried about reaching the maximum number of ECS tasks allowed in the account.

A solutions architect must implement a solution that notifies the development team when Fargate usage reaches 80% of the quota.

What should the architect do?

- A. Use CloudWatch to monitor the Sample Count for each service. Alert when usage exceeds 80%.
- B. Use CloudWatch to monitor ECS service quotas under the AWS/Usage namespace. Create an alarm when utilization exceeds 80%. Notify via SNS.
- C. Use a Lambda function to poll Fargate metrics. Notify via SES when usage exceeds 80%.
- D. Use AWS Config to monitor Fargate quotas. Notify via SES if non-compliant.

Answer: B

Explanation:

B is correct because AWS/Usage namespace in CloudWatch provides built-in metrics for service quota usage. You can create an alarm for Fargate task count usage and notify the dev team using SNS when it reaches 80% of the quota.

A misuses "Sample Count" and doesn't relate to service quota.

C is overly complex and not needed — native metrics exist.

D misuses Config (meant for compliance, not utilization).

Reference:<https://docs.aws.amazon.com/servicequotas/latest/userguide/monitoring-using-cloudwatch.html>

Question: 480

A company needs to migrate some Oracle databases to AWS while keeping others on-premises for compliance. The on-prem databases contain spatial data and run cron jobs. The solution must allow querying on-prem data as foreign tables from AWS.

- A. Use DynamoDB, SCT, and Lambda. Move spatial data to S3 and query with Athena.
- B. Use RDS for SQL Server and AWS Glue crawlers for Oracle access.
- C. Use EC2-hosted Oracle with Application Migration Service. Use Step Functions for cron.
- D. Use RDS for PostgreSQL with DMS and SCT. Use PostgreSQL foreign data wrappers. Connect via Direct Connect.

Answer: D

Explanation:

D is correct because RDS for PostgreSQL supports foreign data wrappers (FDW) that allow querying remote Oracle databases. With AWS Schema Conversion Tool (SCT) and Database Migration Service (DMS), schema and data can be migrated effectively. AWS Direct Connect ensures secure, private connectivity to on-prem databases. Cron jobs can be run via EventBridge or external orchestration. A doesn't support relational/spatial querying.

B doesn't support FDW or spatial types.

C introduces unnecessary complexity.

Reference:<https://www.postgresql.org/docs/current/postgres-fdw.html>

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.PostgreSQL.html

Question: 481

A company hosts an ecommerce site using EC2, ALB, and DynamoDB in one AWS Region. The site uses a custom domain in Route 53. The company wants to replicate the stack to a second Region for disaster recovery and faster access for global customers.

What should the architect do?

- A. Use CloudFormation to deploy to the second Region. Use Route 53 latency-based routing. Enable global tables in DynamoDB.
- B. Use the console to recreate the infra manually in the second Region. Use weighted routing.
- C. Replicate only the S3 and DynamoDB data. Use Route 53 failover routing.
- D. Use Beanstalk and DynamoDB Streams for replication. Use latency-based routing.

Answer: A

Explanation:

A is correct because:

CloudFormation templates enable repeatable infrastructure deployment.

Route 53 latency-based routing ensures users hit the closest Region.

DynamoDB global tables allow multi-Region, active-active replication of application data.

Manual console work (B) is not scalable.

C lacks EC2/ALB replication.

D adds unnecessary services like Beanstalk and doesn't scale cleanly.

Reference: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html>

Question: 482

A company is replicating an application in a secondary Region. The application uses DynamoDB and RDS for MySQL. The secondary Region must function independently during a disaster.

- A. Use DynamoDB global tables and an RDS read replica.
- B. Use DAX and a read replica.
- C. Use global tables and RDS Multi-AZ with standby in secondary Region.
- D. Use Streams and Lambda to copy data. Use read replica.

Answer: A

Explanation:

A is correct because:

DynamoDB global tables allow for multi-Region, active-active usage.

RDS MySQL read replica in another Region supports read workloads and can be promoted during disaster to act as a standalone DB.

B is incorrect: DAX is a cache, not a replication mechanism.

C is wrong because Multi-AZ doesn't span Regions.

D is more manual and error-prone.

Reference: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html>

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

Question: 483

A company runs an application on Amazon EC2 and AWS Lambda

a. The application stores temporary data in Amazon S3. The S3 objects are deleted after 24 hours. The company deploys new versions of the application by launching AWS CloudFormation stacks. The stacks create the required resources. After validating a new version, the company deletes the old stack. The deletion of an old development stack recently failed.

A solutions architect needs to resolve this issue without major architecture changes.

Which solution will meet these requirements?

- A. Create a Lambda function to delete objects from the S3 bucket. Add the Lambda function as a custom resource in the CloudFormation stack with a DependsOn attribute that points to the S3 bucket resource.
- B. Modify the CloudFormation stack to attach a DeletionPolicy attribute with a value of Delete to the S3 bucket.
- C. Update the CloudFormation stack to add a DeletionPolicy attribute with a value of Snapshot for the S3 bucket resource.
- D. Update the CloudFormation template to create an Amazon EFS file system to store temporary files instead of Amazon S3. Configure the Lambda functions to run in the same VPC as the EFS file system.

Answer: A

Explanation:

CloudFormation cannot delete non-empty S3 buckets. Option A allows you to create a custom Lambda resource that deletes all objects in the S3 bucket before the stack deletes it. The DependsOn ensures the bucket deletion occurs only after the Lambda has completed.

B: Adding DeletionPolicy: Delete does not resolve the issue if the bucket still contains objects.

C: Snapshot doesn't apply to S3 and won't help here.

D: Changing to Amazon EFS would require architectural changes, which are not allowed per requirements.

Reference: [https://aws.amazon.com/blogs/devops/safely-delete-s3-buckets-using-aws-](https://aws.amazon.com/blogs/devops/safely-delete-s3-buckets-using-aws-cloudformation/)

[cloudformation/https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html](https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html)

Question: 484

A company runs an application on Amazon EC2 and AWS Lambda

a. The application stores temporary data in Amazon S3. The S3 objects are deleted after 24 hours. The company deploys new versions of the application by launching AWS CloudFormation stacks. The stacks create the required resources. After validating a new version, the company deletes the old stack. The deletion of an old development stack recently failed. A solutions architect needs to resolve this issue without major architecture changes.

Which solution will meet these requirements?

- A. Create a Lambda function to delete objects from an S3 bucket. Add the Lambda function as a custom resource in the CloudFormation stack with a DependsOn attribute that points to the S3 bucket resource.
- B. Modify the CloudFormation stack to attach a DeletionPolicy attribute with a value of Delete to the S3 bucket.
- C. Update the CloudFormation stack to add a DeletionPolicy attribute with a value of Snapshot for the S3 bucket resource.
- D. Update the CloudFormation template to create an Amazon EFS file system to store temporary files instead of Amazon S3. Configure the Lambda functions to run in the same VPC as the EFS file system.

Answer: A

Explanation:

Question: 485

A company has an application that stores user-uploaded videos in an Amazon S3 bucket using S3 Standard storage. Users access videos frequently for the first 180 days, and rarely after that. Most videos are over 100 MB. Users often have poor internet connectivity, and the company uses multipart uploads.

A solutions architect needs to optimize S3 storage costs.

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

- A. Configure the S3 bucket to be a Requester Pays bucket.
- B. Use S3 Transfer Acceleration to upload the videos.
- C. Create a lifecycle rule to expire incomplete multipart uploads after 7 days.
- D. Create a lifecycle rule to transition objects to S3 Glacier Instant Retrieval after 1 day.
- E. Create a lifecycle rule to transition objects to S3 Standard-IA after 180 days.

Answer: C,E

Explanation:

C: Multipart uploads can leave incomplete parts behind, which incur storage costs. Expiring them after 7 days minimizes waste and saves cost.

E: Since objects are infrequently accessed after 180 days, transitioning to S3 Standard-IA is cost effective, especially for large files >128 KB (your 100 MB+ files qualify). IA is for shifting download cost, not reducing your S3 storage expenses.

B helps with upload speed but increases cost.

D is too aggressive; Glacier is not suited for access patterns within the first few days.

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-configuration-examples.html> <https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html>

Question: 486

A company is deploying a new big data analytics cluster across multiple Availability Zones. All nodes must have read/write access to shared file storage that is highly available, POSIX-compatible, and high-throughput.

- A. Use AWS Storage Gateway (file gateway) backed by Amazon S3
- B. Use Amazon EFS in General Purpose performance mode
- C. Use Amazon EBS with Multi-Attach
- D. Use Amazon EFS with Max I/O performance mode

Answer: D

Explanation:

D is the right choice: EFS with Max I/O mode is designed for massively parallel workloads like big data. It supports

high throughput and concurrent access from many EC2 nodes across AZs.
A (Storage Gateway) is not optimized for high-performance, low-latency compute clusters.
B is limited in performance (latency optimized for single-threaded workloads).
C (EBS) cannot be shared across AZs and is not POSIX-compatible for concurrent multi-node access.
Reference: <https://docs.aws.amazon.com/efs/latest/ug/performance.html>

Question: 487

A company is migrating a large on-prem Oracle database (with stored procedures) to AWS. The solution must use managed services, be highly available, and enable a fast migration with minimal downtime.

- A. Use AWS DMS to replicate data to RDS for Oracle. Store database files in S3.
- B. Use backup and restore into EC2-hosted Oracle cluster.
- C. Use DMS to move data to DynamoDB. Recreate stored procedures in Lambda.
- D. Use DMS to migrate to Amazon Aurora PostgreSQL. Use AWS SCT to convert stored procedures.

Answer: D

Explanation:

Dis correct: Aurora PostgreSQL is a fully managed, high-availability database engine.

Use AWS DMS to migrate data and AWS Schema Conversion Tool (SCT) to convert stored procedures from Oracle to Aurora-compatible code.

This meets the requirements of minimal downtime, managed services, and stored procedure support. A is not valid—RDS doesn't allow custom S3 integration for storage.

B involves EC2, which violates the “managed services” constraint.

C requires complete rewrite of the application logic, which contradicts the minimal downtime requirement.

Reference: https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Aurora.html | <https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/Welcome.html>

Question: 488

A company is migrating its on-premises file transfer solution to AWS Transfer Family. The current system includes an SFTP server, a transformation application, and a messaging server. Transformations run every 5 minutes and notify the messaging server when complete.

The company wants to simplify and reduce operational overhead.

- A. Use Amazon EFS and a cron job to perform the transformations. Notify using SNS.
- B. Use Amazon EMR to perform the transformations and notify via SNS.
- C. Use Amazon S3 as storage with AWS Glue triggered by S3 events for transformations, and notify via SQS.
- D. Use Amazon EFS with a time-based AWS Glue job every 5 minutes.

Answer: C

Explanation: C is correct: AWS Transfer Family can store uploaded files directly into Amazon S3. S3 event notifications can trigger AWS Glue jobs to transform the data. Upon successful transformation, Glue can send a message to Amazon SQS, enabling event-driven architecture with minimal management. A and D involve cron jobs or scheduled Glue jobs, which increase operational overhead and delay. B (EMR) is overkill for frequent, lightweight transformations.

Reference: AWS Transfer Family with S3

Question: 489

A company provisions short-lived AWS accounts for students. Each account needs access to ml.p2.xlarge SageMaker instances for training and inference. The default quotas are insufficient. How should quota increases be automated during account provisioning?

- A. Create a quota request template in us-east-1, enable template association, and add quotas for ml.p2.xlarge training and endpoint usage in ap-southeast-2.
- B. Use ml.p2.xlarge training warm pool quota in ap-southeast-2.
- C. Create the template in ap-southeast-2 for SageMaker quotas in us-east-1.
- D. Use warm pool quotas in us-east-1.

Answer: A

Explanation:

A is correct: Quota request templates must be created in us-east-1, the only region that supports them. You specify the desired quota in the correct target Region (ap-southeast-2) for SageMaker training and endpoint usage. This enables automatic quota increases for newly created accounts in the org.

B, C, and D misconfigure either region or quota type.

Reference: Quota request templates

Question: 490

A company runs production workloads on EC2 On-Demand Instances and RDS for PostgreSQL. They want to reduce costs without compromising availability or capacity.

- A. Use CUR and Lambda to terminate underutilized instances. Buy Savings Plans.
- B. Use Budgets and Trusted Advisor, then manually terminate and buy RIs.
- C. Use Compute Optimizer and Trusted Advisor for recommendations. Apply rightsizing, auto scaling, and purchase a Compute Savings Plan.
- D. Use Cost Explorer, alerts, and replace with Spot Instances.

Answer: C

Explanation:

C is correct: AWS Compute Optimizer uses machine learning to analyze usage patterns and recommends rightsizing. Trusted Advisor adds further insights. Combining these with Savings Plans gives the best cost optimization without reducing availability.

A is risky due to using Lambda for termination.

B and D offer partial or manual solutions.

Reference: AWS Compute Optimizer

Question: 491

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 company optimize cost?

- A. Expire S3 objects after 30 days.
- B. Transition S3 content to Glacier Deep Archive after 30 days.
- C. Use Spot Instances with App Runner.
- D. Add auto scaling to Aurora read replica.
- E. Use CloudFront Price Class 200 (Europe & U.S. only).

Answer: B, E

Explanation:

B: Archiving inactive user uploads to Glacier Deep Archive offers the lowest storage cost.

E: Since the users are in Europe only, using Price Class 200 reduces CloudFront delivery cost without affecting performance.

A (expiration) deletes data — not suitable.

C is invalid (App Runner does not support Spot).

D might save little unless read traffic is high.

Reference: CloudFront Price Classes

Question: 492

A company runs a Linux app on Amazon EKS using M6i EC2 instances under a Savings Plan that is about to expire. They want to reduce costs after expiration.

- A. Rebuild containers for ARM64 architecture.
- B. Rebuild containers for container compatibility (invalid/unclear).
- C. Migrate EKS nodes to Graviton (e.g., C7g, M7g).

- D. Replace nodes with latestx86_64instances.
- E. Purchase new Savings Plan for Graviton instance family.
- F. Purchase new Savings Plan for x86_64 instances.

Answer: A,C,E

Explanation:

A: To run on Graviton, containers must supportARM64.

C: Graviton-based EC2 instances offer significant cost savings and better price-performance.

E: Once migrated, aSavings Planfor the new instance family ensures additional cost optimization. B is a non-sensical option.

D and F continue with x86, which is more expensive.

Reference:AWS Graviton Instances

Question: 493

How should EC2 instances in AWS synchronize their clocks with an on-premisesatomic clock NTP server, with theleast administrative overhead?

- A. Configure a DHCP options set with the on-prem NTP server.
- B. Use a custom AMI with Amazon Time Sync.
- C. Deploy a 3rd-party NTP server from Marketplace.
- D. Create an IPsec VPN tunnel to sync over Direct Connect.

Answer: A

Explanation:

Option Ais the most lightweight and scalable approach. By updating the VPC'sDHCP options set, you automatically configure all EC2 instances to use your on-premises NTP server. No extra software or infrastructure is required.

VPC DHCP Options Docs

Question: 494

How can applications in multiple AWS accounts privately access aPostgreSQL RDS instancein a separate AWS account, while managing the number of connections?

- A. Transit Gateway + NAT Gateway
- B. RDS Proxy + PrivateLink via NLB
- C. VPC Peering + Application Load Balancer

D. VPC Peering + NAT Gateway

Answer: B

Explanation:

RDS Proxy optimizes and pools DB connections. By combining this with PrivateLink + NLB, other accounts can securely access the proxy endpoint without going over the internet or setting up peering for every new account.

RDS Proxy Docs

Question: 495

How should a company efficiently process infrequently uploaded S3 data using a long-running (up to 25 minutes) custom application?

- A. ECS on Fargate triggered by EventBridge
- B. Lambda in Step Functions with 30-min timeout
- C. ECS with EC2 and Glue crawler
- D. Lambda triggered by fan-out HTTP EventBridge logic

Answer: A

Explanation:

Amazon ECS on Fargate is ideal for event-driven, long-running jobs with minimal management. Combine S3 event notifications with EventBridge rules to trigger a Fargate task per upload. Using Fargate with EventBridge

Question: 496

How can a company patch EC2 instances without internet access, using a patch source in another account, while accessing Systems Manager and S3?

- A. Custom VPN servers
- B. Transit Gateway + private VIFs
- C. VPC endpoints + VPC peering with patch source
- D. Network ACLs + Transit Gateway

Answer: C

Explanation:

The most secure and manageable solution is to use VPC interface endpoints for S3 and Systems Manager, and establish VPC peering to privately access the patch repo in the core account.

Systems Manager VPC setup

Question: 497

A company is migrating a containerized Kubernetes app with manifest files to AWS. What is the easiest migration path?

- A. App Runner + open-source repo
- B. Amazon EKS with managed node groups and Aurora
- C. ECS on EC2 + task definitions
- D. Rebuild Kubernetes cluster on EC2 manually

Answer: B

Explanation:

Since the company is already using Kubernetes manifests, Amazon EKS is a natural fit. It's a managed Kubernetes control plane and works with Amazon Aurora PostgreSQL, a highly available DB service. **What is EKS?**

Question: 498

A company runs an ecommerce web application on AWS. The static website is hosted on Amazon S3 and served via Amazon CloudFront. API Gateway invokes AWS Lambda for order processing, and Lambda stores data in an Amazon RDS for MySQL DB cluster (On-Demand Instances).

Recently, SQL injection attacks and latency during peak times (cold starts) have been reported. The company wants to ensure scalability, protect against web exploits, and reduce database costs.

- A. Increase Lambda timeout, use RDS Reserved Instances, and use AWS Shield Advanced
- B. Increase Lambda memory, switch to Redshift, use Amazon Inspector
- C. Use provisioned concurrency, switch to Aurora Serverless, use AWS Shield Advanced
- D. Use provisioned concurrency, use RDS Reserved Instances, use AWS WAF with CloudFront

Answer: D

Explanation:

Provisioned concurrency addresses cold starts and improves performance during spikes. RDS Reserved Instance save money over On-Demand for predictable usage.

AWS WAF is designed to protect against SQL injection and web exploits, and integrates with CloudFront.

AWS Lambda Provisioned Concurrency

AWS WAF Web ACL with CloudFront

Question: 499

A VPC spans three Availability Zones, each with public and private subnets. One NAT gateway and one internet gateway exist. Private EC2 instances must connect to the internet.

- A. Add two more NAT gateways (one per AZ). Configure each private subnet to use its AZ's NAT gateway.
- B. Add two more NAT gateways and configure public subnets.
- C. Add internet gateways per AZ and route private subnets.
- D. Add internet gateways per AZ and configure public subnets.

Answer: A

Explanation:

To improve resiliency and fault tolerance, place a NAT Gateway in each AZ, and route each private subnet to its local NAT gateway. This ensures local AZ failure won't break internet access. VPC NAT Gateway High Availability

Question: 500

A company is migrating to AWS and needs to inventory physical and virtual servers, apps, and database relationships to properly rightsize and plan migration.

- A. Use Migration Evaluator with Agentless Collector.
- B. Use Migration Hub with Discovery Agent and Strategy Recommendations.
- C. Use Migration Hub with Agentless Collector and Migration Service.
- D. Use Migration Hub import tool.

Answer: B

Explanation:

Migration Hub Discovery Agent collects deep host-level data, including processes, connections, dependencies, etc. Combined with Strategy Recommendations, it offers accurate insights for planning cloud migration. Migration Hub Strategy Recommendations

Question: 501

An EC2-based ticketing service pulls a frequently updated pricing file (stored in S3) on startup. Sometimes EC2s have stale pricing, causing charge issues.

- A. Lambda updates DynamoDB with new prices.
- B. Lambda updates Amazon EFS.

- C. Use Mountpoint for S3 to mount the pricing file to EC2.
- D. Use Multi-Attach EBS volume for price file.

Answer: C

Explanation:

Mountpoint for Amazon S3 allows EC2 instances to directly access files in S3 as a POSIX-compliant mount point, ensuring they always get the latest data without copying or syncing. It's simple and cost-effective for read-heavy patterns.

Mountpoint for Amazon S3

Question: 502

IoT sensors are manufactured with certificates from a private CA.

- A. They must only connect to AWS after physical installation.
- A. Use Lambda as a pre-provisioning hook to validate serial number before registration.
- B. Use Step Functions to validate before provisioning.
- C. Use Lambda hook but register CA and enable auto-registration.
- D. Use provisioning template and claim certificates without validation.

Answer: A

Explanation:

AWS IoT Core's pre-provisioning hook allows custom logic (like serial number validation) before a device is registered. This ensures that only installed devices connect. Lambda is ideal for such logic. **AWS IoT Just-in-Time Provisioning**

Question: 503

A company creates an Amazon API Gateway API and shares the API with an external development team. The API uses AWS Lambda functions and is deployed to a stage that is named Production. The external development team is the sole consumer of the API. The API experiences sudden increases of usage at specific times, leading to concerns about increased costs. The company needs to limit cost and usage without reworking the Lambda functions.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure the API to send requests to Amazon SQS queues instead of directly to the Lambda functions. Update the Lambda functions to consume messages from the queues and to process the requests. Set up the queues to invoke the Lambda functions when new messages arrive.
- B. Configure provisioned concurrency for each Lambda function. Use AWS Application Auto Scaling to register the Lambda functions as targets. Set up scaling schedules to increase and decrease capacity to match changes in API usage.

- C. Create an API Gateway API key and an AWS WAF Regional web ACL. Associate the web ACL with the Production stage. Add a rate-based rule to the web ACL. In the rule, specify the rate limit and a custom request aggregation that uses the X-API-Key header. Share the API key with the external development team.
- D. Create an API Gateway API key and usage plan. Define throttling limits and quotas in the usage plan. Associate the usage plan with the Production stage and the API key. Share the API key with the external development team.

Answer: D

Explanation:

Question: 504

A company is migrating internal business applications to Amazon EC2 and Amazon RDS in a VPC. The migration requires connecting the cloud-based applications to the on-premises internal network. The company wants to set up an AWS Site-to-Site VPN connection. The company has created two separate customer gateways. The gateways are configured for static routing and have been assigned distinct public IP addresses. Which solution will meet these requirements?

- A. Create one virtual private gateway. Associate the virtual private gateway with the VPC. Enable route propagation for the virtual private gateway in all VPC route tables. Create two Site-to-Site VPN connections with two tunnels for each connection. Configure the Site-to-Site VPN connections to use the virtual private gateway and to use separate customer gateways.
- B. Create one customer gateway. Associate the customer gateway with the VPC. Enable route propagation for the customer gateway in all VPC route tables. Create two Site-to-Site VPN connections with two tunnels for each connection. Configure the Site-to-Site VPN connections to use the customer gateway.
- C. Create two virtual private gateways. Associate the virtual private gateways with the VPC. Enable route propagation for both customer gateways in all VPC route tables. Create two Site-to-Site VPN connections with two tunnels for each connection. Configure the Site-to-Site VPN connections to use separate virtual private gateways and separate customer gateways.
- D. Create two virtual private gateways. Associate the virtual private gateways with the VPC. Enable route propagation for both customer gateways in all VPC route tables. Create four Site-to-Site VPN connections with one tunnel for each connection. Configure the Site-to-Site VPN connections into groups of two. Configure each group to connect to separate customer gateways and separate virtual private gateways.

Answer: A

Explanation:

Question: 505 A company has several Amazon DynamoDB tables in an AWS Region. Each table has more than 100,000 records and was created with default table settings.

To reduce costs, the company needs to identify unused tables. However, the company must maintain the availability and current performance capability of the tables in case the company must use the tables in the future.

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

- A. In Amazon CloudWatch, graph the sum of the ReadThrottleEvents metric and the sum of the WriteThrottleEvents metric for each table over a period of 1 month.
- B. In Amazon CloudWatch, graph the sum of the ConsumedReadCapacityUnits metric and the sum of the ConsumedWriteCapacityUnits metric for each table over a period of 1 month.
- C. Change the provisioned RCUs to 1 for the unused tables. Change the provisioned WCUs to 1 for the unused tables.
- D. Change the capacity mode of the unused tables to on-demand mode.
- E. Change the table class of the unused tables to DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA).
- F. Purchase a reserved capacity of 1 RCU and 1 WCU for each unused table.

Answer: B,D,E

Explanation:

Question: 506

A global ecommerce company has many data centers around the world. With the growth of its stored data, the company needs to set up a solution to provide scalable storage for legacy on-premises file applications. The company must be able to take point-in-time copies of volumes by using AWS Backup and must retain low-latency access to frequently accessed data.

a. The company also needs to have storage volumes that can be mounted as Internet Small Computer System Interface (iSCSI) devices from the company's on-premises application servers. Which solution will meet these requirements?

- A. Provision an AWS Storage Gateway tape gateway. Configure the tape gateway to store data in an Amazon S3 bucket. Deploy AWS Backup to take point-in-time copies of the volumes.
- B. Provision an Amazon FSx File Gateway and an Amazon S3 File Gateway. Deploy AWS Backup to take point-in-time copies of the data.
- C. Provision an AWS Storage Gateway volume gateway in cache mode. Back up the on-premises Storage Gateway volumes with AWS Backup.
- D. Provision an AWS Storage Gateway file gateway in cache mode. Deploy AWS Backup to take point-in-time copies of the volumes.

Answer: C

Explanation:

Question: 507 A company needs to optimize the cost of its application on AWS. The application uses AWS Lambda functions and Amazon ECS containers that run on AWS Fargate. The application is write-heavy and stores data in an Amazon Aurora MySQL database.

The load on the application is not consistent. The application experiences long periods of no usage, followed by sudden and significant increases and decreases in traffic. The database runs on a memory optimized DB instance and has high utilization during peak times. A solutions architect must design a solution that can

scale to handle the changes in traffic.

Which solution will meet these requirements MOST cost-effectively?

- A. Add additional read replicas to the database. Purchase Instance Savings Plans and reserved DB instances for Aurora.
- B. Migrate the database to an Aurora DB cluster that has multiple writer instances. Purchase Instance Savings Plans.
- C. Migrate the database to an Aurora global database. Purchase Compute Savings Plans and reserved DB instances for Aurora.
- D. Migrate the database to Aurora Serverless v2. Purchase Compute Savings Plans.

Answer: D

Explanation:

Question: 508

A company runs an ecommerce website on Amazon ECS behind an Application Load Balancer (ALB). The company stores the container images in Amazon ECR. The website stores data in an Amazon Aurora MySQL DB cluster. The company uses an Amazon S3 bucket to store backup data. The company needs to prevent data tampering. The website domain is registered with Amazon Route 53. The company wants to recreate the setup in a second AWS Region with an RPO of 5 minutes and an RTO of 15 minutes. The company has created an ALB in the second Region. Which solution will meet these requirements?

- A. Create a new ECS deployment that uses the Fargate launch type. Use the ECR repository in the current Region to store and pull container images. Set up a cross-Region read replica in Amazon RDS. Create a backup vault in compliance mode and a backup plan in AWS Backup. Set up a Route 53 primary record in the main Region and a secondary record with a multivalued answer routing policy.
- B. Create a new ECS deployment that uses the Fargate launch type. Use the ECR repository in the current Region to store and pull container images. Set up a cross-Region read replica in Amazon RDS. Set up a Route 53 primary record in the main Region and a secondary record with a failover routing policy.
- C. Set up ECR cross-Region replication. Create a new ECS deployment that uses the Fargate launch type. Migrate the DB cluster to an Aurora global database. Create a backup vault in compliance mode and a backup plan in AWS Backup. Enable point-in-time recovery and cross-Region replication for Amazon S3. Set up a Route 53 primary record in the main Region and a secondary record with a failover routing policy.
- D. Set up ECR cross-Region replication. Create a new ECS deployment that uses the Fargate launch type. Migrate the DB cluster to an Aurora global database. Create a backup vault in governance mode and a backup plan in AWS Backup. Set up a Route 53 primary record in the main Region and a secondary record with a geolocation routing policy.

Answer: C

Explanation:

Question: 509

A company is developing a solution to analyze images. The solution uses a 50 TB reference dataset and analyzes images up to 1 TB in size. The solution spreads requests across an Auto Scaling group of Amazon EC2 Linux instances in a VPC. The EC2 instances are attached to shared Amazon EBS io2 volumes in each Availability Zone. The EBS volumes store the reference dataset.

During testing, multiple parallel analyses led to numerous disk errors, which caused job failures. The company wants the solution to provide seamless data reading for all instances.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a new EBS volume for each EC2 instance. Copy the data from the shared volume to the new EBS volume regularly. Update the application to reference the new EBS volume.
- B. Move all the reference data to an Amazon S3 bucket. Install Mountpoint for Amazon S3 on the EC2 instances. Create gateway endpoints for Amazon S3 in the VPC. Replace the EBS mount point with the S3 mount point.
- C. Move all the reference data to an Amazon S3 bucket. Create an Amazon S3 backed Multi-AZ Amazon EFS volume. Mount the EFS volume on the EC2 instances. Replace the EBS mount point with the EFS mount point.
- D. Upgrade the instances to local storage. Copy the data from the shared EBS volume to the local storage regularly. Update the application to reference the local storage.

Answer: C

Explanation:

Question: 510

A company has deployed applications to thousands of Amazon EC2 instances in an AWS account. A security audit discovers that several unencrypted 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 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:

Question: 511

A company needs to aggregate Amazon CloudWatch logs from its AWS accounts into one central logging account. The collected logs must remain in the AWS Region of creation. The central logging account will then process the logs, normalize the logs into standard output format, and stream the output logs to a security tool for more processing.

A solutions architect must design a solution that can handle a large volume of logging data that needs to be ingested. Less logging will occur outside normal business hours than during normal business hours. The logging solution must scale with the anticipated load. The solutions architect has decided to use an AWS Control Tower design to handle the multi-account logging process. Which combination of steps should the solutions architect take to meet the requirements? (Select THREE.)

- A. Create a destination Amazon Kinesis data stream in the central logging account.
- B. Create a destination Amazon SQS queue in the central logging account.
- C. Create an IAM role that grants Amazon CloudWatch Logs the permission to add data to the Amazon Kinesis data stream. Create a trust policy. Specify the trust policy in the IAM role. In each member account, create a subscription filter for each log group to send data to the Kinesis data stream.
- D. Create an IAM role that grants Amazon CloudWatch Logs the permission to add data to the Amazon SQS queue. Create a trust policy. Specify the trust policy in the IAM role. In each member account, create a single subscription filter for all log groups to send data to the SQS queue.
- E. Create an AWS Lambda function. Program the Lambda function to normalize the logs in the central logging account and to write the logs to the security tool.
- F. Create an AWS Lambda function. Program the Lambda function to normalize the logs in the member accounts and to write the logs to the security tool.

Answer: A,C,E

Explanation:

Question: 512

A company stores application data in many Amazon S3 buckets in one AWS account. Some of the S3 buckets contain sensitive data

- a. The company does not have data inventory for the S3 buckets. The company uses server-side encryption with Amazon S3 managed keys (SSE-S3) to encrypt all data in the S3 buckets.

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

access.

Which solution will meet these requirements?

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

Answer: C

Explanation:

Question: 513

A company is building an application on Amazon EMR to analyze data.

a. The following user groups need to perform different actions:

- Administrator: Provision EMR clusters from different configurations.
- Data engineer: Create an EMR cluster from a small set of available configurations. Run ETL scripts to process data.
- Data analyst: Create an EMR cluster with a specific configuration. Run SQL queries and Apache Hive queries on the data.

A solutions architect must design a solution that gives each group the ability to launch only its authorized EMR configurations. The solution must provide the groups with least privilege access to only the resources that they need. The solution also must provide tagging for all resources that the groups create.

Which solution will meet these requirements?

- A. Configure AWS Service Catalog to control the Amazon EMR versions available for deployment, the cluster configurations, and the permissions for each user group.
- B. Configure Kerberos-based authentication for EMR clusters when the EMR clusters launch. Specify a Kerberos security configuration and cluster-specific Kerberos options.
- C. Create IAM roles for each user group. Attach policies to the roles to define allowed actions for users. Create an AWS Config rule to check for noncompliant resources. Configure the rule to notify the company to address noncompliant resources.
- D. Use AWS CloudFormation to launch EMR clusters with attached resource policies. Create an AWS Config rule to check for noncompliant resources. Configure the rule to notify the company to address noncompliant resources.

Answer: A

Explanation:

Question: 514

A company runs an application on AWS. The application uses an Amazon Aurora MySQL database that is encrypted with the default AWS managed AWS KMS key.

The company must implement a solution to rotate the database encryption key every 180 days. The solution must provide a notification if the encryption key is noncompliant with this standard. Which solution will meet these requirements?

- A. Configure the rotation period for the existing AWS managed KMS key to be 180 days. Implement the cmk-backing-key-rotation-enabled AWS Config managed rule for the existing KMS key. Configure AWS Config to use Amazon SNS to notify the security team if key rotation is noncompliant.
- B. Create a new AWS managed KMS key with automatic rotation set for 180 days. Take a snapshot of the database. Restore the snapshot to a new Aurora cluster that uses the new KMS key. Create an AWS Config custom rule that uses an AWS Lambda function to validate the key rotation period. Configure AWS Config to use Amazon SES to notify the security team if key encryption is noncompliant.
- C. Create a new customer managed KMS key with automatic rotation set for 180 days. Take a snapshot of the database. Restore the snapshot to a new Aurora cluster that uses the new KMS key. Create an AWS Config custom rule that uses an AWS Lambda function to validate the key rotation period. Configure AWS Config to use Amazon SNS to notify the security team if key encryption is noncompliant.
- D. Create a new customer managed KMS key with automatic rotation set for 180 days. Update the database to use the new KMS key for encryption. Implement the cmk-backing-key-rotation-enabled AWS Config managed rule for the new KMS key. Configure AWS Config to use Amazon SES to notify the security team if key rotation is noncompliant.

Answer: B

Explanation:

Question: 515 A company uses an organization in AWS Organizations that has multiple AWS accounts. The accounts host multiple resources that are tagged with a CostCenter tag key. The tag value is the name of the team. The company wants to accurately identify the cost of the resources so that the company can charge each team accordingly.

Which solution meets these requirements?

- A. Activate the CostCenter user-defined tag in the organization's management account. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Use the tag breakdown in the report to obtain the total cost for the resources that have the CostCenter tag.
- B. Activate the CostCenter user-defined tag in every member account. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Create an AWS Lambda function that runs monthly to retrieve the reports and calculate the total cost for the resources that have the CostCenter tag.
- C. Activate the CostCenter user-defined tag in every member account. Schedule a monthly AWS Cost and

Usage Report from the management account. Use the tag breakdown in the report to calculate the total cost for the resources that have the CostCenter tag.

D. Customize a report in the AWS Trusted Advisor organization view. Configure the report to generate monthly billing summaries for resources that have the CostCenter tag under the AWS accounts.

Answer: A

Explanation:

Question: 516

A company is running a three-tier web application in an on-premises data center. The frontend is a PHP application that is served by an Apache web server. The middle tier is a monolithic Java SE application. The storage tier is a 60 TB PostgreSQL database.

The three-tier web application recently crashed and became unresponsive. The database also reached capacity because of read operations. The company wants to migrate to AWS to resolve these issues and improve scalability,

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

- A. Configure an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer to host the web server. Use Amazon EFS for the frontend static assets.
- B. Host the static single-page application on Amazon S3. Use an Amazon CloudFront distribution to serve the application.
- C. Create a Docker container to run the Java SE application. Use AWS Fargate to host the container.
- D. Create an AWS Elastic Beanstalk environment for Java to host the Java SE application.
- E. Migrate the PostgreSQL database to an Amazon EC2 instance that is larger than the on-premises PostgreSQL database.
- F. Use AWS DMS to replatform the PostgreSQL database to an Amazon Aurora PostgreSQL database. Use Aurora Auto Scaling for read replicas.

Answer: A,D,F

Explanation:

Question: 517

A company is planning to migrate workloads from its on-premises data center to Amazon EC2 instances. The workloads run on physical servers and VMware virtual servers. The company has gathered details about each on-premises server and virtual server, including server specification, CPU utilization, and memory utilization.

The company has stored these details in a .csv file named **onprem.csv**.

Before the migration, the company must estimate the cost of running the servers on AWS and must determine recommended EC2 instance types for the servers. The company must export this information to a different .csv file.

Which solution will meet these requirements?

- www.atmicnetworks.com
- A. Configure AWS Compute Optimizer to generate recommendations from an external source. Import the onprem.csv file. Export the Compute Optimizer recommendations to a new .csv file.
 - B. Import the onprem.csv file into AWS Migration Hub by using AWS Migration Hub import. Use EC2 instance recommendations from Migration Hub to generate recommendations. Export the recommendations to a new .csv file.
 - C. Deploy AWS Application Discovery Service Agentless Collector on premises. Use Agentless Collector to import the onprem.csv file. Send the file to AWS Migration Hub. Use EC2 instance recommendations from Migration Hub to generate recommendations. Export the recommendations to a new .csv file.
 - D. Upload the onprem.csv file to an Amazon S3 bucket. Configure Migration Evaluator to import the data from the S3 bucket. Generate and confirm recommendations by using Migration Evaluator Quick Insights. Export the final recommendations to a new .csv file in the S3 bucket.

Answer: D

Explanation:

Question: 518

A company wants to retire its Oracle Solaris NFS storage arrays. The company requires rapid data migration over its internet network connection to a combination of destinations for Amazon S3, Amazon EFS, and Amazon FSx for Windows File Server. The company also requires a full initial copy, as well as incremental transfers of changes until the retirement of the storage arrays. All data must be encrypted and checked for integrity. What should a solutions architect recommend to meet these requirements?

- A. Configure AWS Application Migration Service. Create a project and deploy the AWS Replication Agent and token to the storage array. Run the migration plan to start the transfer.
- B. Configure AWS DataSync. Configure the DataSync agent and deploy it to the local network. Create a transfer task and start the transfer.
- C. Configure the aws S3 sync command. Configure the AWS client on the client side with credentials. Run the sync command to start the transfer.
- D. Configure AWS Transfer for FTP. Configure the FTP client with credentials. Script the client to connect and sync to start the transfer.

Answer: B

Explanation:

Question: 519

A company uses AWS Organizations. The company creates a central VPC in an AWS account that is designated for networking in a single AWS Region. The central VPC has an AWS Site-to-Site VPN connection to the company's on-premises network. A solutions architect must create another AWS account that uses the

same networking resources that the central VPC uses.

Which solution meets these requirements MOST cost-effectively?

- A. Create a VPC in the new AWS account. Create a new Site-to-Site VPN connection for the onpremises connection.
- B. Use AWS Resource Access Manager to share the VPN connection in the central VPC with the new AWS account.
- C. Create a VPC in the new AWS account. Configure a virtual private gateway to connect to the central VPC.
- D. Use AWS Resource Access Manager to share the subnets in the central VPC with the new AWS account.

Answer: D

Explanation:

Question: 520

A company uses a software package for surveys. During surveys, data is uploaded from a field operator's device to an Amazon S3 bucket. A custom application that runs on several Amazon EC2 instances polls the S3 bucket for new data.

a. When new data is available, the software processes the data.

The data uploads are infrequent. The processing software can take up to 25 minutes to analyze each data upload. The company wants to optimize the application workflow to process the S3 data. Which solution will meet these requirements with the LEAST operational overhead?

A. Modify the application to accept new S3 object keys as inputs. Containerize the application. Deploy the container to an Amazon ECS cluster that uses the AWS Fargate launch type. Configure S3 bucket notifications to send events to Amazon EventBridge when new objects are uploaded. Create an EventBridge rule that invokes an ECS task to run the application when a new S3 object event OCCURS.

B. Modify the application to accept new S3 object keys as inputs. Containerize the application. Deploy the container image to AWS Lambda functions. Create a new AWS Step Functions state machine to invoke the Lambda functions. Configure the state machine with a Task state that calls the Lambda functions. Set the Task state's Timeout property to 30 minutes.

C. Modify the application to accept new S3 object keys as inputs. Move the application from EC2 instances to Amazon ECS by using the EC2 capacity provider. Create an AWS Glue crawler to check the S3 bucket and invoke the application. Configure the application to process the data when the data is uploaded to Amazon S3.

D. Modify the application to use HTTP to poll new S3 object keys that reference data to process. Containerize the application. Deploy the container image to AWS Lambda functions. Configure S3 bucket notifications to send events to Amazon EventBridge when new objects are uploaded. Create an EventBridge rule that invokes the Lambda functions to post the new objects to HTTP endpoints by using fan-out.

Answer: A

Explanation:

Question: 521

A company is planning a large event where a promotional offer will be introduced. The company's website is hosted on AWS and backed by an Amazon RDS for PostgreSQL DB instance. The website explains the promotion and includes a sign-up page that collects user information and preferences. Management expects large and unpredictable volumes of traffic periodically, which will create many database writes. A solutions architect needs to build a solution that does not change the underlying data model and ensures that submissions are not dropped before they are committed to the database.

Which solutions meets these requirements?

- A. Immediately before the event, scale up the existing DB instance to meet the anticipated demand. Then scale down after the event.
- B. Use Amazon SQS to decouple the application and database layers. Configure an AWS Lambda function to write items from the queue into the database.
- C. Migrate to Amazon DynamoDB and manage throughput capacity with automatic scaling.
- D. Use Amazon ElastiCache (Memcached) to increase write capacity to the DB instance.

Answer: B

Explanation:

Question: 522

A company has many AWS accounts in an organization in AWS Organizations. The accounts contain many Amazon EC2 instances that run different types of workloads. The workloads have different usage patterns.

The company needs recommendations for how to rightsize the EC2 instances based on CPU and memory usage during the last 90 days.

Which combination of steps will provide these recommendations? (Select THREE.)

- A. Opt in to AWS Compute Optimizer and enable trusted access for Compute Optimizer for the organization.
- B. Configure a delegated administrator account for AWS Systems Manager for the organization.
- C. Use an AWS CloudFormation stack set to enable detailed monitoring for all the EC2 instances.
- D. Install and configure the Amazon CloudWatch agent on all the EC2 instances to send memory utilization metrics to CloudWatch.
- E. Activate enhanced metrics in AWS Compute Optimizer.
- F. Configure AWS Systems Manager to pass metrics to AWS Trusted Advisor.

Answer: A,D,F

Explanation:

Question: 523

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

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

Which solution meets these requirements?

- A. Provision an Application Load Balancer (ALB) in front of the game servers. Create an Amazon CloudFront distribution that has no geographical restrictions. Set the ALB as the origin. Perform DNS lookups for the cloudfront.net domain name. Use the resulting IP addresses in the game's client application.
- B. Provision game servers in each AWS Region. Provision an Application Load Balancer in front of the game servers. Create an Amazon Route 53 latency-based routing policy for the game's client application to use with DNS lookups.
- C. Provision game servers in each AWS Region. Provision a Network Load Balancer (NLB) in front of the game servers. Create an accelerator in AWS Global Accelerator, and configure endpoint groups in each Region. Associate the NLBs with the corresponding Regional endpoint groups. Point the game client's application to the Global Accelerator endpoints.
- D. Provision game servers in each AWS Region. Provision a Network Load Balancer (NLB) in front of the game servers. Create an Amazon CloudFront distribution that has no geographical restrictions. Set the NLB as the origin. Perform DNS lookups for the cloudfront.net domain name. Use the resulting IP addresses in the game's client application.

Answer: C

Explanation:

Question: 524

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

Which solution will meet these requirements?

- A. Create an AWS CloudTrail trail in each account. Specify CloudTrail management events for the trail. Configure CloudTrail to send the events to Amazon CloudWatch Logs. Configure CloudWatch crossaccount observability. Query the data in CloudWatch Logs Insights.
- B. Use a delegated administrator account to create an AWS CloudTrail Lake data store. Specify CloudTrail management events for the data store. Enable the data store for all accounts in the organization. Query the data in CloudTrail Lake.
- C. Use a delegated administrator account to create an AWS CloudTrail trail. Specify CloudTrail management events for the trail. Enable the trail for all accounts in the organization. Keep all other settings as default. Query the CloudTrail data from the CloudTrail event history page.
- D. Use AWS CloudFormation StackSets to deploy AWS CloudTrail Lake data stores in each account. Specify CloudTrail management events for the data stores. Keep all other settings as default. Query the data in

CloudTrail Lake.

Answer: B

Explanation:

Question: 525

A company has a sales system that stores transactions as .csv files in an Amazon S3 bucket. The S3 bucket is configured to use S3 Intelligent-Tiering. Most of the .csv files are between 64 KB and 100 KB in size. All rows and columns of the .csv files must be read when the data is processed. The company must keep the data for 5 years.

The company stores several million xsv files every day. The company must minimize the cost of storing and querying the xsv files.

Which solution will meet these requirements?

- A. Create an AWS Glue job to convert the .csv files into Apache Parquet format. Use Amazon S3 to invoke the AWS Glue job every time a .csv file arrives.
- B. Create an AWS Glue job to compress the .csv files. Schedule the AWS Glue job every hour to compress the files for the previous hour into one .csv file.
- C. Create an AWS Lambda function to convert the .csv files into Apache Parquet format. Use Amazon S3 to invoke the Lambda function every time a .csv file arrives.
- D. Create an AWS Lambda function to compress the .csv files. Use Amazon S3 to invoke the Lambda function every time a .csv file arrives.

Answer: A

Explanation:

Question: 526

A company has 20 accounts in an organization in AWS Organizations. The accounts are in two OUs: development and production. Multiple teams use the development accounts.

The company wants to control the cost that is associated with the development accounts. The company needs a solution that provides a notification when the forecasted monthly cost for all development accounts exceeds a threshold.

A solutions architect creates an Amazon SNS topic and subscribes an email address to the topic. What should the solutions architect do next to meet the notification requirement with the LEAST configuration effort?

- A. Enable Amazon CloudWatch billing alerts in the organization's management account. Create a CloudWatch billing alarm by configuring the EstimatedCharges metric for each development account as a linked account. Configure the SNS topic for email alerts when the EstimatedCharges metric value exceeds the threshold.
- B. Create an AWS Cost and Usage Report in the organization's management account. Configure report delivery to an Amazon S3 bucket. Configure an AWS Glue job to extract the report data into Amazon Athena. Configure

AWS Step Functions to analyze the consolidated cost of all the development accounts. Configure the SNS topic for email alerts when the cost exceeds the threshold.

C. Use AWS Budgets to create a cost budget in the organization's management account. Configure each development account as a linked account. Configure an alert threshold. Configure the SNS topic for email alerts.

D. Enable AWS Cost Explorer in the organization's management account. Configure each development account as a linked account. Configure an alert threshold. Configure the SNS topic for email alerts.

Answer: C

Explanation:

Question: 527

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:

Question: 528

A company runs a web application on a single Amazon EC2 instance. End users experience slow application performance during times of peak usage, when CPU utilization is consistently more than 95%.

A user data script installs required custom packages on the EC2 instance. The process of launching the instance takes several minutes.

The company is creating an Auto Scaling group that has mixed instance groups, varied CPUs, and a maximum

capacity limit. The Auto Scaling group will use a launch template for various configuration options. The company needs to decrease application latency when new instances are launched during auto scaling. Which solution will meet these requirements?

- A. Use a predictive scaling policy. Use an instance maintenance policy to run the user data script. Set the default instance warmup time to 0 seconds.
- B. Use a dynamic scaling policy. Use lifecycle hooks to run the user data script. Set the default instance warmup time to 0 seconds.
- C. Use a predictive scaling policy. Enable warm pools for the Auto Scaling group. Use an instance maintenance policy to run the user data script.
- D. Use a dynamic scaling policy. Enable warm pools for the Auto Scaling group. Use lifecycle hooks to run the user data script.

Answer: D

Explanation:

Question: 529

A company has a website that serves many visitors. The company deploys a backend service for the website in a primary AWS Region and a disaster recovery (DR) Region.

A single Amazon CloudFront distribution is deployed for the website. The company creates an Amazon Route 53 record set with health checks and a failover routing policy for the primary Region's backend service. The company configures the Route 53 record set as an origin for the CloudFront distribution. The company configures another record set that points to the backend service's endpoint in the DR Region as a secondary failover record type. The TTL for both record sets is 60 seconds.

Currently, failover takes more than 1 minute. A solutions architect must design a solution that will provide the fastest failover time.

Which solution will achieve this goal?

- A. Deploy an additional CloudFront distribution. Create a new Route 53 failover record set with health checks for both CloudFront distributions.
- B. Set the TTL to 1 second for the existing Route 53 record sets that are used for the backend service in each Region.
- C. Create new record sets for the backend services by using a latency routing policy. Use the record sets as an origin in the CloudFront distribution.
- D. Create a CloudFront origin group that includes two origins, one for each backend service Region. Configure origin failover as a cache behavior for the CloudFront distribution.

Answer: D

Explanation:

Question: 530

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 THREE.)

- A. Use a Network Load Balancer (NLB) in front of the player-matching instance. Use a friendly DNS entry in Amazon Route 53-point address.
- B. Use an Application Load Balancer (ALB) in front of the player-matching instance. Use a friendly DNS entry in Amazon Route 53 p 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 balance

Answer: A,D

Explanation:

Question: 531

A company needs to migrate its on-premises database fleet to Amazon RDS. The company is currently using a mixture of Microsoft SQL Server, and Oracle databases. Some of the databases have custom schemas and stored procedures.

Which combination of steps should the company take for the migration? (Select TWO.)

- A. Use Migration Evaluator Quick Insights to analyze the source databases and to identify the stored procedures that need to be migrated.
- B. Use AWS Application Migration Service to analyze the source databases and to identify the stored procedures that need to be migrated.
- C. Use AWS SCT to analyze the source databases for changes that are required.
- D. Use AWS DM5 to migrate the source databases to Amazon RD5.
- E. Use AWS DataSync to migrate the data from the source databases to Amazon RDS.

Answer: C

Explanation:

Question: 532

A company has an application that uses Amazon EC2 instances in an Auto Scaling group. The quality assurance (QA) department needs to launch and test the application. The application environments are currently launched by the manager of the department using an AWS CloudFormation template. To launch the stack, the manager

uses a role with permission to use CloudFormation, EC2, and Auto Scaling APIs. The manager wants to allow QA to launch environments, but does not want to grant broad permissions to each user.

Which set up would achieve these goals?

- A. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to assume the manager's role, restricts the permissions to the template and the resources it creates. Train users to launch the template from the CloudFormation console.
- B. Create an AWS Service Catalog product from the environment template. Add a launch constraint to the product with the existing manager's department permission to use AWS Service Catalog APIs only. Train users to launch the template from the AWS Service Catalog console.
- C. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to use CloudFormation and restrict the permissions to the template and the resources it creates. Train users to launch the template from the CloudFormation console.
- D. Create an AWS Elastic Beanstalk application from the environment template. Give users in the QA department permission to use Elastic Beanstalk only. Train users to launch Elastic Beanstalk environments with the Elastic Beanstalk CLI, passing the existing role to the environment.

Answer: B

Explanation:

Question: 533

A company migrated to AWS and uses AWS Business Support. The company wants to monitor the cost-effectiveness of Amazon EC2 instances. 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 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 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 a 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 a QuickSight dashboard that uses the DynamoDB table as a data source to identify and stop underutilized development EC2 instances.

Answer: B

Explanation:

Question: 534

A company has deployed production workloads on Amazon EC2 On-Demand Instances and Amazon RDS for PostgreSQL DB instances in multiple environments. The company has the AWS Business Support plan. A solutions architect must optimize the cost of the workloads without negatively affecting the availability or compute capacity of the workloads. Which solution will meet these requirements?

- A. Use AWS Cost and Usage Reports to analyze the most expensive instances and usage patterns. Use AWS Lambda to terminate underutilized instances. Purchase Compute Savings Plans for instances for highly utilized workloads.
- B. Use AWS Budgets to track spending for each environment. Configure AWS Trusted Advisor cost optimization checks to rightsize instances. Create billing alerts in Amazon CloudWatch. Terminate underutilized instances. Purchase Reserved Instances for highly utilized workloads.
- C. Opt in to AWS Compute Optimizer. Use Compute Optimizer and AWS Trusted Advisor to identify underutilized instances. Implement recommendations from Compute Optimizer, modify instance types, rightsize instances, and apply Auto Scaling groups. Purchase a Compute Savings Plan.
- D. Use AWS Cost Explorer recommendations to rightsize underutilized instances. Create billing alerts in Amazon CloudWatch. Replace the EC2 On-Demand Instances with Spot Instances for underutilized instances. Stop any instances that are not in use.

Answer: C

Explanation:

Comprehensive and Detailed

Option C is the most comprehensive and aligns with AWS best practices for cost optimization without compromising availability or performance:

AWS Compute Optimizer provides recommendations to rightsize EC2 instances, EBS volumes, and Lambda functions based on utilization metrics. By opting in, the company can receive tailored suggestions for optimizing resource configurations.

AWS Trusted Advisor offers cost optimization checks, including identifying idle or underutilized resources, which can be rightsized or terminated to reduce costs.

Implementing recommendations from Compute Optimizer and Trusted Advisor ensures that resources are appropriately sized, leading to cost savings without affecting performance.

Auto Scaling groups help maintain application availability and allow automatic scaling of EC2 instances based on demand, ensuring that compute capacity aligns with workload requirements. Compute Savings Plans provide flexible pricing for EC2, AWS Lambda, and AWS Fargate usage, offering significant savings over On-Demand pricing in exchange for a commitment to a consistent amount of usage.

This approach ensures cost optimization while maintaining the necessary compute capacity and availability for production workloads.

Reference:

AWS Compute Optimizer: Provides recommendations for optimizing AWS resources based on utilization metrics.

AWS Trusted Advisor: Offers real-time guidance to help provision resources following AWS best practices.

AWS Auto Scaling: Automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost.

AWS Savings Plans: Flexible pricing model offering lower prices compared to On-Demand pricing, in exchange for a commitment to a consistent amount of usage.

Question: 535

A company is using a single AWS Region for its ecommerce website. The website includes a web application that runs on several Amazon EC2 instances behind an Application Load Balancer (ALB). The website also includes an Amazon DynamoDB table. A custom domain name in Amazon Route 53 is linked to the ALB. The company created an SSL/TLS certificate in AWS Certificate Manager (ACM) and attached the certificate to the ALB. The company is not using a content delivery network as part of its design. The company wants to replicate its entire application stack in a second Region to provide disaster recovery, plan for future growth, and provide improved access time to users. A solutions architect needs to implement a solution that achieves these goals and minimizes administrative overhead. Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Create an AWS CloudFormation template for the current infrastructure design. Use parameters for important system values, including Region. Use the CloudFormation template to create the new infrastructure in the second Region.
- B. Use the AWS Management Console to document the existing infrastructure design in the first Region and to create the new infrastructure in the second Region.
- C. Update the Route 53 hosted zone record for the application to use weighted routing. Send 50% of the traffic to the ALB in each Region.
- D. Update the Route 53 hosted zone record for the application to use latency-based routing. Send traffic to the ALB in each Region.
- E. Update the configuration of the existing DynamoDB table by enabling DynamoDB Streams. Add the second Region to create a global table.
- F. Create a new DynamoDB table. Enable DynamoDB Streams for the new table. Add the second Region to create a global table. Copy the data from the existing DynamoDB table to the new table as a one-time operation.

Answer: A, D, E

Explanation:

Comprehensive and Detailed

A . Using AWS CloudFormation templates allows for the automation of infrastructure deployment. By parameterizing the templates, resources can be provisioned consistently across multiple Regions, reducing administrative overhead and ensuring infrastructure as code practices.

D . Amazon Route 53 latency-based routing directs user requests to the AWS Region that provides the lowest latency, improving access times for users globally and enhancing the user experience. E . Enabling DynamoDB Streams and adding a second Region to create a global table ensures that data is replicated across Regions, providing high availability and disaster recovery capabilities. This combination of steps ensures that the application stack is replicated efficiently across Regions, providing disaster recovery, improved performance, and scalability, all while minimizing manual administrative tasks.

Reference:

AWS CloudFormation: Enables modeling and setting up AWS resources so that you can spend less time managing those resources and more time focusing on your applications.

Amazon Route 53 Latency-Based Routing: Helps improve performance for your global applications by serving user requests from the AWS Region that provides the lowest latency.

Amazon DynamoDB Global Tables: Provides a fully managed, multi-Region, and multi-active database that delivers fast and local read and write performance for massively scaled, global applications.

Question: 536

A company has dozens of AWS accounts for different teams, applications, and environments. The company has defined a custom set of controls that all accounts must have. The company is concerned that potential misconfigurations in the accounts could lead to security issues or noncompliance. A solutions architect must design a solution that deploys the custom controls by using infrastructure as code (IaC) in a repeatable way. Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure AWS Config rules in each account to evaluate the account settings against the custom controls. Define AWS Lambda functions in AWS CloudFormation templates. Program the Lambda functions to remediate noncompliant AWS Config rules. Deploy the CloudFormation templates as stack sets during account creation. Configure the stack sets to invoke the Lambda functions.
- B. Configure AWS Systems Manager associations to remediate configuration issues across accounts. Define the desired configuration state in an AWS CloudFormation template by using AWS::SSM::Association. Deploy the CloudFormation templates as stack sets to all accounts during account creation.
- C. Enable AWS Control Tower to set up and govern the multi-account environment. Use blueprints that enforce security best practices. Use Customizations for AWS Control Tower and CloudFormation templates to define the custom controls for each account. Use Amazon EventBridge to deploy Customizations for AWS Control Tower during account-provisioning lifecycle events.
- D. Enable AWS Security Hub in all the accounts to aggregate findings in a central administrator account. Develop AWS CloudFormation templates to create Amazon EventBridge rules, AWS Lambda functions, and CloudFormation stacks in each account to remediate Security Hub findings. Deploy the CloudFormation stacks during account provisioning to set up the automated remediation.

Answer: C

Explanation:

Comprehensive and Detailed

Option C offers a scalable and low-overhead solution for managing custom controls across multiple AWS accounts:

AWS Control Tower provides a pre-configured environment to set up and govern a secure, multiaccount AWS environment based on AWS best practices.

Customizations for AWS Control Tower (CfCT) allows for the deployment of custom configurations and resources, such as AWS Config rules and IAM policies, across accounts and organizational units using AWS CloudFormation templates.

Amazon EventBridge integrates with AWS Control Tower to automate the deployment of customizations during account provisioning events, ensuring that all new accounts adhere to the defined controls without manual intervention.

This approach ensures consistent enforcement of custom controls across all accounts with minimal

operational overhead.

Reference:

AWS Control Tower: Automates the setup of a baseline environment, or landing zone, that is a secure, well-architected multi-account AWS environment.

Customizations for AWS Control Tower: Enables you to customize your AWS Control Tower landing zone using AWS CloudFormation templates and service control policies (SCPs).

Amazon EventBridge: A serverless event bus that makes it easier to build event-driven applications at scale using events generated from your applications, integrated SaaS applications, and AWS services.

Question: 537

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

- A. Create a network ACL that blocks outbound traffic on port 80. Associate the network ACL with all subnets in the application account. In the application account and the core account, deploy one EC2 instance that runs a custom VPN server. Create a VPN tunnel to access the private VPC. Update the route table in the application account.
- B. Create private VIFs for Systems Manager and Amazon S3. Delete the NAT gateway from the VPC in the application account. Create a transit gateway to access the patch source repository EC2 instances in the core account. Update the route table in the core account.
- C. Create VPC endpoints for Systems Manager and Amazon S3. Delete the NAT gateway from the VPC in the application account. Create a VPC peering connection to access the patch source repository EC2 instances in the core account. Update the route tables in both accounts.
- D. Create a network ACL that blocks inbound traffic on port 80. Associate the network ACL with all subnets in the application account. Create a transit gateway to access the patch source repository EC2 instances in the core account. Update the route tables in both accounts.

Answer: C

Explanation:

Option C is the correct and most efficient solution, aligning with AWS best practices for secure and private connectivity:

Create VPC Endpoints for Systems Manager and Amazon S3:

Systems Manager VPC Endpoints: By creating interface VPC endpoints for Systems Manager

(com.amazonaws.region.ssm, com.amazonaws.region.ec2messages, and

com.amazonaws.region.ssmmessages), the EC2 instances can communicate with Systems Manager services

without requiring internet access. This setup ensures that patching operations can be conducted securely

within the AWS network.

Amazon S3 VPC Endpoint: A gateway VPC endpoint for Amazon S3 (com.amazonaws.region.s3) allows EC2 instances to access S3 buckets privately. This is essential for accessing application data stored in S3 without traversing the public internet.

Reference: docs.aws.amazon.com

Delete the NAT Gateway:

Removing the NAT gateway ensures that EC2 instances in the application account cannot access the internet, satisfying the requirement to prevent internet access. This action enhances the security posture by eliminating a potential vector for unauthorized outbound traffic.

Create a VPC Peering Connection:

Establishing a VPC peering connection between the application account's VPC and the core account's private VPC enables direct, private communication between the EC2 instances in both accounts. This setup allows the application account's EC2 instances to access the patch source repository hosted in the core account securely.

Reference: docs.aws.amazon.com

Update Route Tables in Both Accounts:

After setting up the VPC peering connection, it's crucial to update the route tables in both VPCs to allow traffic to flow between them. This configuration ensures that the EC2 instances in the application account can reach the patch source repository in the core account and vice versa.

Why Other Options Are Incorrect:

Option A: Implementing a custom VPN solution introduces unnecessary complexity and operational overhead. Additionally, merely blocking outbound traffic on port 80 does not comprehensively prevent internet access, as other ports (e.g., 443 for HTTPS) remain open.

Option B: Creating private virtual interfaces (VIFs) is typically associated with AWS Direct Connect, which is not applicable in this scenario. Moreover, using a transit gateway, while feasible, is more complex and may be unnecessary for this use case.

Option D: Blocking inbound traffic on port 80 does not prevent outbound internet access. Furthermore, employing a transit gateway adds complexity and cost, which may not be justified given the requirements.

Conclusion:

Option C provides a secure, efficient, and cost-effective solution that meets all the specified requirements:
Prevents EC2 instances from accessing the internet.

Enables access to Amazon S3 and Systems Manager services via VPC endpoints.

Facilitates secure communication with the patch source repository in the core account through VPC peering.

This approach leverages AWS's native networking features to maintain a secure and private environment for patch management operations.

Question: 538

A company runs an ecommerce web application on AWS. The web application is hosted as a static website on Amazon S3 with Amazon CloudFront for content delivery. An Amazon API Gateway API invokes AWS Lambda functions to handle user requests and order processing for the web application. The Lambda functions store data in an Amazon RDS for MySQL DB cluster that uses On-Demand Instances. The DB cluster usage has been consistent in the past 12 months. Recently, the website has experienced SQL injection and web exploit attempts. Customers also report that order processing time has increased during periods of peak usage. During these periods, the Lambda functions often have cold starts. As the company grows, the company needs to ensure scalability and low-latency access during traffic peaks. The company also must optimize the database

costs and add protection against the SQL injection and web exploit attempts. Which solution will meet these requirements?

- A. Configure the Lambda functions to have an increased timeout value during peak periods. Use RDS Reserved Instances for the database. Use CloudFront and subscribe to AWS Shield Advanced to protect against the SQL injection and web exploit attempts.
- B. Increase the memory of the Lambda functions. Transition to Amazon Redshift for the database. Integrate Amazon Inspector with CloudFront to protect against the SQL injection and web exploit attempts.
- C. Use Lambda functions with provisioned concurrency for compute during peak periods. Transition to Amazon Aurora Serverless for the database. Use CloudFront and subscribe to AWS Shield Advanced to protect against the SQL injection and web exploit attempts.
- D. Use Lambda functions with provisioned concurrency for compute during peak periods. Use RDS Reserved Instances for the database. Integrate AWS WAF with CloudFront to protect against the SQL injection and web exploit attempts.

Answer: D

Explanation:

Option D best addresses all aspects of the problem:

Provisioned Concurrency for Lambda ensures that the Lambda functions have pre-initialized execution environments ready, eliminating cold start delays and ensuring low-latency performance during traffic spikes.

RDS Reserved Instances provide significant cost savings for workloads with consistent and predictable usage, which aligns with the company's usage patterns.

AWS WAF integrated with CloudFront directly addresses the security concern by providing application layer protection against SQL injection and web exploit attempts.

This solution ensures scalability, performance consistency, cost savings, and security protection without introducing unnecessary complexity.

Question: 539

A company needs to move some on-premises Oracle databases to AWS. The company has chosen to keep some of the databases on premises for business compliance reasons. The on-premises databases contain spatial data and run cron jobs for maintenance. The company needs to connect to the on-premises systems directly from AWS to query data as a foreign table. Which solution will meet these requirements?

- A. Create Amazon DynamoDB global tables with auto scaling enabled. Use AWS SCT and AWS DMS to move the data from on premises to DynamoDB. Create an AWS Lambda function to move the spatial data to Amazon S3. Query the data by using Amazon Athena. Use Amazon EventBridge to schedule jobs in DynamoDB for maintenance. Use Amazon API Gateway for foreign table support.
- B. Create an Amazon RDS for Microsoft SQL Server DB instance. Use native replication to move the data from on premises to the DB instance. Use AWS SCT to modify the SQL Server schema as needed after replication. Move the spatial data to Amazon Redshift. Use stored procedures for system maintenance. Create AWS Glue crawlers to connect to the on-premises Oracle databases for foreign table support.
- C. Launch Amazon EC2 instances to host the Oracle databases. Place the EC2 instances in an Auto Scaling group. Use AWS Application Migration Service to move the data from on premises to the EC2 instances and for real-time bidirectional change data capture (CDC) synchronization. Use Oracle native spatial data support.

Create an AWS Lambda function to run maintenance jobs as part of an AWS Step Functions workflow. Create an internet gateway for foreign table support.

D. Create an Amazon RDS for PostgreSQL DB instance. Use AWS SCT and AWS DMS to move the data from on premises to the DB instance. Use PostgreSQL native spatial data support. Run cron jobs on the DB instance for maintenance. Use AWS Direct Connect to connect the DB instance to the on-premises environment for foreign table support.

Answer: D

Explanation:

Option D is the most appropriate solution as it addresses all the specified requirements:

Amazon RDS for PostgreSQL supports native spatial data types through the PostGIS extension, making it suitable for handling spatial data migrated from Oracle databases.

AWS Schema Conversion Tool (SCT) and AWS Database Migration Service (DMS) can be used to migrate data from on-premises Oracle databases to Amazon RDS for PostgreSQL, facilitating the transition while preserving data integrity.

Cron jobs can be scheduled directly on the PostgreSQL DB instance to handle maintenance tasks, aligning with the company's existing maintenance practices.

AWS Direct Connect provides a dedicated network connection between the on-premises environment and AWS, enabling secure and efficient querying of on-premises databases as foreign tables from the RDS instance.

This solution ensures compatibility with spatial data, maintains existing maintenance workflows, and provides a secure and efficient connection to on-premises systems.

Question: 540

A company uses AWS Organizations with all features enabled to manage its accounts. The company has configured AWS Backup to run every 4 hours on several Amazon EFS mount points in the eu-west-2 Region. The backups are stored in the default vault. The company needs a disaster recovery (DR) plan that restores into the eu-west-1 Region and a specific recovery account. The backups must be encrypted at all times. Which solution will meet these requirements?

- A. Configure AWS Resource Access Manager (AWS RAM) to share the backup vault with the recovery account. Create a new backup vault in the recovery account. Encrypt the data by using an AWS managed KMS key. Schedule a copy job in the recovery account to copy the backup vault to the new vault.
- B. Create a new backup vault in the source account and a new backup vault in the recovery account. Encrypt the data by using a multi-Region customer managed KMS key. Redirect the backups to the new backup vault. Configure a key policy statement to allow access to the key from the recovery account. Schedule a cross-account backup plan to the recovery account.
- C. Create an Amazon S3 bucket. Create a new multi-Region customer managed KMS key to encrypt the S3 bucket data. Schedule a copy job from the backup vault that copies the data to the S3 bucket. Configure cross-account access for the recovery account to the S3 bucket. Schedule a second copy job in the recovery account to copy the data from the S3 bucket into the default vault.
- D. Configure AWS DataSync to copy the EFS data to eu-west-1 in the source account. In the recovery account,

create a new backup vault. Encrypt the data by using an AWS managed KMS key. In the source account, schedule a cross-account backup plan to the recovery account's vault in eu-west-1.

Answer: B

Explanation:

Option B effectively addresses the disaster recovery requirements:

Creating a new backup vault in both the source and recovery accounts allows for organized and secure storage of backups.

Using a multi-Region customer managed KMS key ensures that the backups are encrypted at all times and can be decrypted in the recovery account, maintaining data security during cross-account and cross-Region transfers.

Redirecting backups to the new backup vault in the source account and configuring a key policy to allow access from the recovery account facilitates seamless and secure backup copying.

Scheduling a cross-account backup plan enables automated and regular copying of backups to the recovery account, ensuring that the DR plan is consistently maintained.

This approach ensures that backups are securely and efficiently replicated to the recovery account in a different Region, aligning with the company's DR requirements.

Question: 541

A company is replicating an application in a secondary AWS Region. The application in the primary Region reads from and writes to several Amazon DynamoDB tables. The application also reads customer data from an Amazon RDS for MySQL DB instance. The company plans to use the secondary Region as part of a disaster recovery plan. The application in the secondary Region must function without dependencies on the primary Region. Which solution will meet these requirements with the LEAST development effort?

A. Configure DynamoDB global tables. Replicate the required tables to the secondary Region. Create a read replica of the RDS DB instance in the secondary Region. Configure the secondary application to use the DynamoDB tables and the read replica in the secondary Region.

B. Use DynamoDB Accelerator (DAX) to cache the required tables in the secondary Region. Create a read replica of the RDS DB instance in the secondary Region. Configure the secondary application to use DAX and the read replica in the secondary Region.

C. Configure DynamoDB global tables. Replicate the required tables to the secondary Region. Enable Multi-AZ for the RDS DB instance. Configure the standby replica to be created in the secondary Region. Configure the secondary application to use the DynamoDB tables and the standby replica in the secondary Region.

D. Set up DynamoDB streams from the primary Region. Process the streams in the secondary Region to populate new DynamoDB tables. Create a read replica of the RDS DB instance in the secondary Region. Configure the secondary application to use the DynamoDB tables and the read replica in the secondary Region.

Answer: A

Explanation:

Option A provides a straightforward and efficient solution:

DynamoDB global tables automatically replicate data across multiple Regions, ensuring that the secondary Region has up-to-date data without the need for custom replication logic.

Creating a read replica of the RDS DB instance in the secondary Region allows the application to access customer data without relying on the primary Region.

Configuring the secondary application to use these resources ensures that it can function independently, fulfilling the disaster recovery requirements with minimal development effort. This solution leverages AWS's managed services to provide a resilient and low-maintenance disaster recovery setup.

Question: 542

A global company has a mobile app that displays ticket barcodes. Customers use the tickets on the mobile app to attend live events. Event scanners read the ticket barcodes and call a backend API to validate the barcode data against data in a database. After the barcode is scanned, the backend logic writes to the database's single table to mark the barcode as used. The company needs to deploy the app on AWS with a DNS name of `api.example.com`. The company will host the database in three AWS Regions around the world. Which solution will meet these requirements with the LOWEST latency?

- A. Host the database on Amazon Aurora global database clusters. Host the backend on three Amazon ECS clusters that are in the same Regions as the database. Create an accelerator in AWS Global Accelerator to route requests to the nearest ECS cluster. Create an Amazon Route 53 record that maps `api.example.com` to the accelerator endpoint.
- B. Host the database on Amazon Aurora global database clusters. Host the backend on three Amazon EKS clusters that are in the same Regions as the database. Create an Amazon CloudFront distribution with the three clusters as origins. Route requests to the nearest EKS cluster. Create an Amazon Route 53 record that maps `api.example.com` to the CloudFront distribution.
- C. Host the database on Amazon DynamoDB global tables. Create an Amazon CloudFront distribution. Associate the CloudFront distribution with a CloudFront function that contains the backend logic to validate the barcodes. Create an Amazon Route 53 record that maps `api.example.com` to the CloudFront distribution.
- D. Host the database on Amazon DynamoDB global tables. Create an Amazon CloudFront distribution. Associate the CloudFront distribution with a Lambda@Edge function that contains the backend logic to validate the barcodes. Create an Amazon Route 53 record that maps `api.example.com` to the CloudFront distribution.

Answer: D

Explanation:

Option D offers the lowest latency and highest efficiency:

Amazon DynamoDB global tables replicate data across multiple Regions, ensuring that the database is available and responsive in each Region.

Lambda@Edge allows the backend logic to execute at AWS edge locations, reducing the distance between the user's device and the code execution environment, thereby minimizing latency. Amazon CloudFront serves as the content delivery network, routing requests to the nearest edge location where the Lambda@Edge function is executed.

Amazon Route 53 maps the DNS name to the CloudFront distribution, ensuring that users are directed to the

optimal edge location.

This architecture ensures rapid response times and high availability for a globally distributed user base.

Question: 543

A company's web application uses an Amazon API Gateway API, AWS Lambda functions, and Amazon DynamoDB global tables to handle backend requests. The web application is deployed in two AWS Regions in an active-passive model. The company uses Amazon Route 53 for DNS. The web application requires a manual DNS update to fail over to the secondary Region. An analytics Lambda function runs in the same AWS account. The function has caused Lambda concurrency to reach 90% of the current quota on an average day. A recent surge in traffic for the analytics workload resulted in throttled Lambda requests and a poor user experience for the web application users. A solutions architect must increase the reliability of the web application. The solution must use an Amazon CloudWatch alarm to send an Amazon SNS notification when the Lambda concurrency reaches a specific utilization threshold. Which solution will meet these requirements with the LEAST operational overhead?

- A. Set reserved concurrency on the web application Lambda functions. Implement Route 53 health checks and failover records to route traffic to the secondary Region. Configure the CloudWatch alarm to use the AWS Trusted Advisor ServiceLimitUsage metric and to send the SNS notification.
- B. Set reserved concurrency on the web application Lambda functions. Implement Route 53 health checks and latency records to route traffic to the secondary Region. Configure the CloudWatch alarm to use the AWS Trusted Advisor ServiceLimitUsage metric and to send an SNS notification.
- C. Set provisioned concurrency on the web application Lambda functions. Implement Route 53 health checks and failover records to route traffic to the secondary Region. Configure the CloudWatch alarm to use the Lambda ConcurrentExecutions metric and to send an SNS notification.
- D. Set provisioned concurrency on the web application Lambda functions. Implement Route 53 health checks and geolocation records to route traffic to the secondary Region. Configure the CloudWatch alarm to use the Lambda ProvisionedConcurrencyInvocations metric and to send an SNS notification.

Answer: C**Explanation:**

The use of provisioned concurrency ensures the web application's Lambda functions have preinitialized execution environments, removing cold start latency and maintaining performance during high-traffic periods.

Route 53 health checks and failover records automate DNS failover to the secondary Region, improving application availability and reliability.

The CloudWatch alarm is configured to monitor the Lambda ConcurrentExecutions metric and send an SNS notification if concurrency usage nears the limit, enabling quick operational response. This approach minimizes manual management, ensuring reliability and performance during peak traffic while meeting best practices for AWS Lambda and Route 53 failover.

Question: 544

A company in the United States (US) has acquired a company in Europe. Both companies use the AWS Cloud. The US company has built a new application with a microservices architecture. The US company is hosting the application across five VPCs in the us-east-2 Region. The application must be able to access resources in one VPC in the eu-west-1 Region. However, the application must not be able to access any other VPCs. The VPCs in both Regions have no overlapping CIDR ranges. All accounts are already consolidated in one organization in AWS Organizations. Which solution will meet these requirements MOST cost-effectively?

- A. Create one transit gateway in eu-west-1. Attach the VPCs in us-east-2 and the VPC in eu-west-1 to the transit gateway. Create the necessary route entries in each VPC so that the traffic is routed through the transit gateway.
- B. Create one transit gateway in each Region. Attach the involved subnets to the regional transit gateway. Create the necessary route entries in the associated route tables for each subnet so that the traffic is routed through the regional transit gateway. Peer the two transit gateways.
- C. Create a full mesh VPC peering connection configuration between all the VPCs. Create the necessary route entries in each VPC so that the traffic is routed through the VPC peering connection.
- D. Create one VPC peering connection for each VPC in us-east-2 to the VPC in eu-west-1. Create the necessary route entries in each VPC so that the traffic is routed through the VPC peering connection.

Answer: A

Explanation:

Creating a single transit gateway in eu-west-1 allows for centralized and cost-effective connectivity between the VPCs involved, without the complexity of creating multiple peering connections or transit gateways.

This approach also allows for selective attachment of only the required VPCs, meeting the strict security and access requirements.

Route table configuration ensures traffic flows properly through the transit gateway, providing connectivity for the application without exposing other VPC resources.

This solution is the most cost-effective and operationally efficient choice for secure inter-Region communication.

Question: 545

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, create a cost allocation tag that is named BusinessUnit. 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:

The correct approach involves activating the BusinessUnit cost allocation tag in the management account to ensure costs are tagged consistently across the organization.

The AWS Cost and Usage Report (CUR), delivered to a single Amazon S3 bucket in the management account, consolidates cost and usage data across all member accounts.

Amazon Athena can efficiently query the CUR data directly from S3, while Amazon QuickSight can be used to create interactive dashboards and reports that visualize the cost allocation by business unit. This solution minimizes complexity, centralizes reporting, and leverages AWS native services for the best cost optimization and insights.

Question: 546

A company's inventory application stores data in an Amazon Aurora PostgreSQL DB cluster in a single AWS Region. The company wants to improve resiliency by extending the database infrastructure to a secondary Region. The company wants an RTO of 15 minutes and an RPO of 5 minutes. The solution must not run Aurora DB instances in the secondary Region when the application is operational in the primary Region. Which solution meets these requirements?

- A. Configure AWS DMS to copy the Aurora DB cluster in the primary Region to the secondary Region. Use AWS DMS to synchronize the primary DB cluster with the secondary DB cluster.
- B. Create a new Aurora PostgreSQL DB cluster in the secondary Region. Use AWS Backup to synchronize the primary DB cluster with the secondary DB cluster.
- C. Create a headless Aurora DB cluster in the second Region that is part of the same global DB cluster as the primary Region's DB cluster.
- D. Create an AWS Backup job to back up the DB cluster and copy the DB cluster to the secondary Region every 5 minutes.

Answer: C

Explanation:

Creating a headless Aurora DB cluster in the secondary Region as part of an Aurora global database ensures continuous, low-latency replication without running active instances in the secondary Region.

This configuration meets the RTO and RPO requirements by allowing rapid promotion of the secondary cluster in a failover scenario.

It also reduces operational costs compared to continuously running standby instances in the secondary Region.

This aligns with AWS best practices for DR and cross-Region resiliency.

Question: 547

A company is migrating its on-premises file transfer solution to AWS Transfer Family. The on-premises host includes an SFTP server to receive files, an application that performs a transformation of the files, and a messaging server. The transformations run every 5 minutes. When a transformation is complete, the application sends a message to a queue on the messaging server. The company needs to simplify the solution and reduce the management of the components. What should the company do to meet these requirements with the LEAST operational overhead?

- A. Configure Transfer Family to use Amazon EFS storage. Use a cron job on Amazon EFS to perform the transformations. Configure the cron job to publish a message to an Amazon SNS topic when a file has been transformed.
 - B. Configure Transfer Family to use Amazon S3 storage. Use Amazon EMR to perform the transformations. Configure Amazon EMR to send a message to an Amazon SNS topic when a file has been transformed.
 - C. Configure Transfer Family to use Amazon S3 storage. Use AWS Glue to perform the transformations after S3 event notifications. Configure AWS Glue to send a message to an Amazon SQS queue when a file has been transformed.
 - D. Configure Transfer Family to use Amazon EFS storage. Create an AWS Glue time-based job to run every 5 minutes to initiate an AWS Glue transformation. Configure AWS Glue to send a message to an Amazon SQS queue when a file has been transformed.
-

Answer: C

Explanation:

Using AWS Transfer Family with Amazon S3 storage eliminates the need to manage file servers and offers a scalable and durable file storage solution.

S3 event notifications automatically trigger the AWS Glue transformation jobs upon file arrival, eliminating manual job scheduling.

AWS Glue jobs can be configured to send a notification to an Amazon SQS queue after completion, maintaining the messaging workflow with minimal operational overhead.

This fully managed approach simplifies the architecture and leverages AWS-native automation.

Question: 548

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? (Select 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:

Configuring an S3 Lifecycle policy to expire incomplete multipart uploads after 7 days prevents storage of partially uploaded objects, avoiding unnecessary costs from failed uploads. Additionally, implementing a lifecycle policy to transition objects to S3 Standard-IA after 180 days ensures older videos that are rarely accessed move to a lower-cost storage class, significantly reducing storage costs. These measures are aligned with AWS cost optimization best practices for S3 data lifecycle management.

Question: 549

A company has an online learning platform that teaches data science. The platform uses the AWS Cloud to provision on-demand lab environments for its students. Each student receives a dedicated AWS account for a short time. Students need access to ml.p2.xlarge instances to run a single Amazon SageMaker machine learning training job and to deploy the inference endpoint. Account provisioning is automated. The accounts are members of an organization in AWS Organizations with all features enabled. The accounts must be provisioned in the ap-southeast-2 Region. The default resource usage quotas are not sufficient for the accounts. A solutions architect must enhance the account provisioning process to include automated quota increases. Which solution will meet these requirements?

- A. Create a quota request template in the us-east-1 Region in the organization's management account. Enable template association. Add a quota for SageMaker in ap-southeast-2 for ml.p2.xlarge training job usage. Set the desired quota to 1. Add a quota for SageMaker in ap-southeast-2 for ml.p2.xlarge endpoint usage. Set the desired quota to 1.
- B. Create a quota request template in the us-east-1 Region in the organization's management account. Enable template association. Add a quota for SageMaker in ap-southeast-2 for ml.p2.xlarge training warm pool usage. Set the desired quota to 2.
- C. Create a quota request template in ap-southeast-2 in the organization's management account. Enable template association. Add a quota for SageMaker in the us-east-1 Region for ml.p2.xlarge training job usage.

Set the desired quota to 1. Add a quota for SageMaker in us-east-1 for ml.p2.xlarge endpoint usage. Set the desired quota to 1.

D. Create a quota request template in ap-southeast-2 in the organization's management account. Enable template association. Add a quota for SageMaker in the us-east-1 Region for ml.p2.xlarge training warm pool usage. Set the desired quota to 2.

Answer: A

Explanation: Creating a quota request template in the us-east-1 Region of the management account ensures it applies to all newly provisioned accounts in the organization.

By specifying the required ml.p2.xlarge training job and endpoint usage quotas in ap-southeast-2, the accounts will automatically receive these higher quotas upon creation, supporting the student lab environments without additional manual quota requests.

This process integrates seamlessly with AWS Organizations for automated and standardized account provisioning.

Question: 550

A company runs a processing engine in the AWS Cloud. The engine processes environmental data from logistics centers to calculate a sustainability index. The company has millions of devices in logistics centers that are spread across Europe. The devices send information to the processing engine through a RESTful API. The API experiences unpredictable bursts of traffic. The company must implement a solution to process all data that the devices send to the processing engine. Data loss is unacceptable. Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB) for the RESTful API. Create an Amazon SQS queue. Create a listener and a target group for the ALB. Add the SQS queue as the target. Use a container that runs in Amazon ECS with the Fargate launch type to process messages in the queue.
- B. Create an Amazon API Gateway HTTP API that implements the RESTful API. Create an Amazon SQS queue. Create an API Gateway service integration with the SQS queue. Create an AWS Lambda function to process messages in the SQS queue.
- C. Create an Amazon API Gateway REST API that implements the RESTful API. Create a fleet of Amazon EC2 instances in an Auto Scaling group. Create an API Gateway Auto Scaling group proxy integration. Use the EC2 instances to process incoming data.
- D. Create an Amazon CloudFront distribution for the RESTful API. Create a data stream in Amazon Kinesis Data Streams. Set the data stream as the origin for the distribution. Create an AWS Lambda function to consume and process data in the data stream.

Answer: B

Explanation:

Using Amazon API Gateway with an SQS queue ensures all data sent to the API is durably stored for processing, even during traffic surges.

API Gateway service integration with SQS provides a managed, serverless ingestion layer that absorbs

unpredictable traffic bursts without data loss.

The AWS Lambda function processes messages from SQS, providing scalable, on-demand compute to handle environmental data efficiently.

This solution aligns with AWS best practices for building resilient, serverless, event-driven architectures that can handle high-volume, bursty workloads.

Question: 551

A 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 DynamoDB to look up pricing.
- B. Create an AWS Lambda function to update an 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 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:

Using Mountpoint for Amazon S3 to directly mount the S3 bucket to the EC2 instances ensures that the latest version of the pricing file is always available without repeatedly downloading it to each instance.

This eliminates stale data issues and minimizes data transfer costs since the file does not need to be copied each time an instance launches.

This solution is cost-effective and aligns with AWS best practices for accessing frequently updated shared files directly from S3.

Question: 552

A retail company is mounting IoT sensors in all of its stores worldwide. During the manufacturing of each sensor, the company's private certificate authority (CA) issues an X.509 certificate that contains a unique serial number. The company then deploys each certificate to its respective sensor.

A solutions architect needs to give the sensors the ability to send data to AWS after they are installed.

Sensors must not be able to send data to AWS until they are installed.

Which solution will meet these requirements?

A. Create an AWS Lambda function that can validate the serial number. Create an AWS IoT Core provisioning template. Include the SerialNumber parameter in the Parameters section. Add the Lambda function as a pre-provisioning hook. During manufacturing, call the RegisterThing API operation and specify the template and parameters.

B. Create an AWS Step Functions state machine that can validate the serial number. Create an AWS IoT Core provisioning template. Include the SerialNumber parameter in the Parameters section.

Specify the Step Functions state machine to validate parameters. Call the StartThingRegistrationTask API operation during installation.

C. Create an AWS Lambda function that can validate the serial number. Create an AWS IoT Core provisioning template. Include the SerialNumber parameter in the Parameters section. Add the Lambda function as a pre-provisioning hook. Register the CA with AWS IoT Core, specify the provisioning template, and set the allow-auto-registration parameter.

D. Create an AWS IoT Core provisioning template. Include the SerialNumber parameter in the Parameters section. Include parameter validation in the template. Provision a claim certificate and a private key for each device that uses the CA. Grant AWS IoT Core service permissions to update AWS IoT things during provisioning.

Answer: C

Explanation:

By integrating a Lambda pre-provisioning hook into the AWS IoT Core provisioning template, the serial number can be validated against the CA before the device is allowed to send data.

Using allow-auto-registration ensures that only installed devices can register and send data, meeting the security requirement that sensors cannot send data until physically installed.

This approach follows AWS IoT's secure onboarding pattern for managing device identity and authorization.

Question: 553

A company collects air quality data from sensors. The company plans to use the MQTT protocol to send the data to AWS IoT Core. The company will process the data and then will store the data in an Amazon Aurora database.

During periods of low air quality, sensors will send data more frequently. The company must buffer the data during these periods to make sure that no data is lost before the data is processed and stored.

Which solution will meet these requirements?

A. Create an Amazon Kinesis data stream. Create an AWS IoT rule action and set the data stream as the target. Create an AWS Step Functions state machine that is invoked by the data stream. Use the state machine to process and store the data.

B. Create an Amazon Kinesis data stream. Create an AWS IoT rule action and set the data stream as the target. Create an application that runs on an Amazon ECS cluster with the AWS Fargate launch type. Configure the application to read data from the data stream, process the data, and store the data.

C. Create an Amazon SQS queue. Create an AWS IoT rule action and set the SQS queue as the target. Create an AWS Step Functions state machine that is invoked by the SQS queue. Use the state machine to process and store the data.

D. Create an Amazon SNS topic. Create an AWS IoT rule action and set the SNS topic as the target. Create an application that runs on an Amazon ECS cluster with the AWS Fargate launch type. Configure the application to read data from the SNS topic, process the data, and store the data.

Answer: B

Explanation:

Using an Amazon Kinesis data stream as a target for the IoT Core rule ensures durable and scalable buffering of data during traffic spikes.

An ECS Fargate application reads data from Kinesis, processes it, and stores it in Aurora, meeting the requirement for flexible, serverless data handling.

This solution guarantees no data loss while providing real-time or near-real-time data processing capabilities.

Question: 554

A solutions architect has created a single VPC on AWS. The VPC has one internet gateway and one NAT gateway. The VPC extends across three Availability Zones. Each Availability Zone includes one public subnet and one private subnet. The three private subnets contain Amazon EC2 instances that must be able to connect to the internet.

Which solution will increase the network resiliency of this architecture?

A. Add two NAT gateways so that each Availability Zone has a NAT gateway. Configure a route table for each private subnet to send traffic to the NAT gateway in the subnet's Availability Zone.

B. Add two NAT gateways so that each Availability Zone has a NAT gateway. Configure a route table for each public subnet to send traffic to the NAT gateway in the subnet's Availability Zone.

C. Add two internet gateways so that each Availability Zone has an internet gateway. Configure a route table for each private subnet to send traffic to the internet gateway in the subnet's Availability Zone.

D. Add two internet gateways so that each Availability Zone has an internet gateway. Configure a route table for each public subnet to send traffic to the internet gateway in the subnet's Availability Zone.

Answer: A

Explanation:

By adding two more NAT gateways (one per Availability Zone), each private subnet has direct access to a NAT gateway in its own AZ.

This increases network resiliency because even if a single NAT gateway becomes unavailable, the others remain operational, ensuring consistent internet connectivity for the private EC2 instances.

This follows AWS best practices for highly available network architecture in multi-AZ deployments.

Question: 555

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. Ensure that NTP traffic is allowed between AWS and the on-premises networks.
- 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:

Using a DHCP options set in AWS allows EC2 instances to automatically receive network configuration information, including the IP address of the on-premises NTP server.

This ensures that all EC2 instances in the VPC consistently use the on-premises atomic clock as their authoritative time source, maintaining the same time synchronization as the existing environment. By allowing NTP traffic over the Direct Connect connection and ensuring proper security group and network ACL configurations, this approach eliminates the need to manually configure each EC2 instance's NTP settings, significantly reducing administrative overhead and ensuring consistent, accurate time across the environment.

This aligns with AWS best practices for extending on-premises time synchronization to EC2 instances in a hybrid environment.

Question: 556

A company wants to migrate its website to AWS. The website uses microservices and runs on containers that are deployed in an on-premises, self-managed Kubernetes cluster. All the manifests that define the deployments for the containers in the Kubernetes deployment are in source control. All data for the website is stored in a PostgreSQL database. An open source container image repository runs alongside the on-premises environment.

A solutions architect needs to determine the architecture that the company will use for the website on AWS.

Which solution will meet these requirements with the LEAST effort to migrate?

- A. Create an AWS App Runner service. Connect the App Runner service to the open source container image repository. Deploy the manifests from on premises to the App Runner service. Create an Amazon RDS for PostgreSQL database.
- B. Create an Amazon EKS cluster that has managed node groups. Copy the application containers to a new Amazon ECR repository. Deploy the manifests from on premises to the EKS cluster. Create an Amazon Aurora PostgreSQL DB cluster.
- C. Create an Amazon ECS cluster that has an Amazon EC2 capacity pool. Copy the application containers to a new Amazon ECR repository. Register each container image as a new task definition. Configure ECS services for each task definition to match the original Kubernetes deployments. Create an Amazon Aurora PostgreSQL DB cluster.
- D. Rebuild the on-premises Kubernetes cluster by hosting the cluster on Amazon EC2 instances. Migrate the open source container image repository to the EC2 instances. Deploy the manifests from on premises to the new cluster on AWS. Deploy an open source PostgreSQL database on the new cluster.

Answer: B

Explanation:

Migrating to an Amazon EKS cluster with managed node groups minimizes the effort required because: EKS is fully managed, offering native Kubernetes support, making it easy to deploy the existing Kubernetes manifests without major changes.

Copying containers to Amazon ECR allows for fully managed, scalable container image storage in AWS, eliminating reliance on the on-premises container repository.

Deploying the existing manifests directly to EKS reuses all the existing configuration, such as service definitions, deployments, and scaling policies, simplifying migration.

Using Amazon Aurora PostgreSQL provides a fully managed, highly available database service that is compatible with PostgreSQL, reducing operational overhead compared to managing a self-hosted database.

This approach leverages the power of AWS managed services while preserving the existing microservices and deployment practices, ensuring minimal disruption and fastest migration path.

Question: 557

A company has an application that uses an on-premises Oracle database. The company is migrating the database to the AWS Cloud. The database contains customer data and stored procedures. The company needs to migrate the database as quickly as possible with minimum downtime. The solution on AWS must provide high availability and must use managed services for the database. Which solution will meet these requirements?

- A. Use AWS DMS to replicate data from the on-premises Oracle database to a new Amazon RDS for Oracle database. Transfer the database files to an Amazon S3 bucket. Configure the RDS database to use the S3 bucket as database storage. Set up S3 replication for high availability. Redirect the application to the RDS DB instance.
- B. Create a database backup of the on-premises Oracle database. Upload the backup to an Amazon S3 bucket.

Shut down the on-premises Oracle database to avoid any new transactions. Restore the backup to a new Oracle cluster that consists of Amazon EC2 instances across two Availability Zones. Redirect the application to the EC2 instances.

C. Use AWS DMS to replicate data from the on-premises Oracle database to a new Amazon DynamoDB table. Use DynamoDB Accelerator (DAX) and implement global tables for high availability. Rewrite the stored procedures in AWS Lambda. Run the stored procedures in DAX. After replication, redirect the application to the DAX cluster endpoint.

D. Use AWS DMS to replicate data from the on-premises Oracle database to a new Amazon Aurora PostgreSQL database. Use AWS SCT to convert the schema and stored procedures. Redirect the application to the Aurora DB cluster.

Answer: D

Explanation:

Using AWS Database Migration Service (DMS) with an Amazon Aurora PostgreSQL target and the AWS Schema Conversion Tool (SCT) ensures a rapid migration with minimal downtime.

DMS continuously replicates data, supporting near-zero downtime migration while keeping the source and target in sync.

SCT automatically converts the Oracle schema and stored procedures to PostgreSQL-compatible format, minimizing manual effort.

Aurora provides a managed, highly available database that scales across multiple AZs, ensuring high availability for the migrated workload.

This method leverages AWS-managed services for database and migration, ensuring a secure, reliable, and low-overhead transition to the cloud.

Question: 558

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: D

Explanation:

Creating an internal ALB and configuring it as a PrivateLink endpoint service enables private connectivity between internal applications and the metadata API, ensuring that traffic does not traverse the public internet.

Internal ALB: Ensures traffic stays within the AWS network and is not exposed publicly. **PrivateLink endpoint service:** Provides secure, private access to the ALB from the internal EC2 instances in other AWS accounts.

Traffic stays within the AWS global network, leveraging AWS security best practices and meeting the new policy requirements for no public internet exposure.

This approach is secure, scalable, and minimizes management complexity compared to API Gateway solutions.

Question: 559

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:

To minimize EKS compute costs:

A . Rebuild the application container images to support ARM64 architecture: Graviton instances (ARM-based) require container images built for the ARM64 architecture. This ensures compatibility and allows leveraging Graviton-based compute.

C . Migrate the EKS nodes to the most recent generation of Graviton-based instances: AWS Graviton instances (e.g., c7g, m7g) offer significant price/performance benefits over x86_64 instances. By migrating EKS worker nodes to Graviton, the company can reduce compute costs substantially.

E . Purchase a new EC2 Instance Savings Plan for the newly selected Graviton instance family: Savings Plans provide cost savings compared to On-Demand pricing. After moving to Graviton, purchasing a new Savings Plan tailored for these instance families ensures continued cost savings.

This approach aligns with AWS cost optimization best practices—migrating to the latest instance types and purchasing Savings Plans to match the new usage pattern—while ensuring full compatibility with EKS.

Question: 560

A company is planning to migrate to the AWS Cloud. The company hosts many applications on Windows servers and Linux servers. Some of the servers are physical, and some of the servers are virtual. The company uses several types of databases in its on-premises environment. The company does not have an accurate inventory of its on-premises servers and applications.

The company wants to rightsize its resources during migration. A solutions architect needs to obtain information about the network connections and the application relationships. The solutions architect must assess the company's current environment and develop a migration plan.

Which solution will provide the solutions architect with the required information to develop the migration plan?

- A. Use Migration Evaluator to request an evaluation of the environment from AWS. Use the AWS Application Discovery Service Agentless Collector to import the details into a Migration Evaluator Quick Insights report.
- B. Use AWS Migration Hub and install the AWS Application Discovery Agent on the servers. Deploy the Migration Hub Strategy Recommendations application data collector. Generate a report by using Migration Hub Strategy Recommendations.
- C. Use AWS Migration Hub and run the AWS Application Discovery Service Agentless Collector on the servers. Group the servers and databases by using AWS Application Migration Service. Generate a report by using Migration Hub Strategy Recommendations.
- D. Use the AWS Migration Hub import tool to load the details of the company's on-premises environment. Generate a report by using Migration Hub Strategy Recommendations.

Answer: B

Explanation:

To develop a migration plan with accurate inventory and dependency data:

AWS Migration Hub provides a single view for tracking migration tasks and resources across multiple AWS services.

The AWS Application Discovery Agent (installed on servers) collects detailed data about running processes, system performance, and network connections.

Migration Hub Strategy Recommendations leverages this data to automatically identify application patterns, generate recommended AWS target services, and provide a detailed migration plan. This approach ensures accurate data collection, detailed dependency mapping, and tailored recommendations—crucial for a successful and right-sized migration to AWS.

Question: 561

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. Set up an S3 gateway endpoint in that 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: C

Question: 562

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: D, E, F