

Amazon.SAP-C02.vFeb-2024.by.Json.129q

Website: [www.VCEplus.io](http://www.VCEplus.io)

Twitter: [https://twitter.com/VCE\\_Plus](https://twitter.com/VCE_Plus)

**SAP-C02**

**Exam Name: AWS Certified Solutions Architect - Professional**



## Exam A

### QUESTION 1

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.

**Correct Answer: C**

**Section:**

**Explanation:**

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

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

www.VCEplus.io

### QUESTION 2

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 cross-account 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.

**Correct Answer: B**

**Section:**

**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 manually.

### QUESTION 3

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:Reoon 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.

**Correct Answer: C, E**

**Section:**

**Explanation:**

[https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_examples\\_aws\\_deny-requested-region.html](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](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_ec2.html)

### QUESTION 4

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.

**Correct Answer: C**

**Section:**

**Explanation:**

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-orgs-manage-auto-deployment.html>

### QUESTION 5

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.

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.

**Correct Answer: D**

**Section:**

**Explanation:**

<https://docs.aws.amazon.com/waf/latest/developerguide/marketplace-managed-rule-groups.html>

#### QUESTION 6

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.

**Correct Answer: B**

**Section:**

**Explanation:**

<https://docs.aws.amazon.com/firehose/latest/dev/creating-the-stream-to-splunk.html>

#### QUESTION 7

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.

**Correct Answer: A, C, F**

**Section:**

**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 8

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 in 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.

**Correct Answer: A, C**

**Section:**

**Explanation:**

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

#### QUESTION 9

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

Which strategy should the solutions architect use?

- A. Use AWS Lambda to run the application. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours.
- B. Use AWS Batch to run the application. Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours.
- C. Use AWS Fargate to run the application. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours.
- D. Use Amazon EC2 Spot Instances to run the application. Use AWS CodeDeploy to deploy and run the application every 4 hours.

**Correct Answer: C**

**Section:**

**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 10

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

**Correct Answer: A**

**Section:**

**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 11

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

**Correct Answer: D**

**Section:**

**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 credentials. You can use the `sts get-session-token` AWS CLI command to request temporary credentials that include an MFA token. 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 present. 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.

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_temp.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp.html)

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_mfa\\_enable\\_cliapi.html](https://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/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 12

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.

**Correct Answer: C**

**Section:**

**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.

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 13

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.

**Correct Answer: A**

**Section:**

**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 14

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.

**Correct Answer: A, E**

**Section:**

**Explanation:**

A Direct Connect gateway allows you to connect multiple VPCs across different Regions to a Direct Connect connection<sup>1</sup>. A public VIF allows you to access AWS public services such as EC2<sup>1</sup>. 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 VPCs<sup>2</sup>. This solution is cheaper than setting up additional Direct Connect connections or using a private VIF with VPC peering.

#### QUESTION 15

A company ingests and processes streaming market data. 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.

**Correct Answer: B**

**Section:**

**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 storage<sup>1</sup>. 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 Instances<sup>2</sup>. 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 16

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.

**Correct Answer: C**

**Section:**

**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 premises<sup>1</sup>. AWS Outposts allows customers to run some AWS services locally and connect to a broad range of services available in the local AWS Region<sup>1</sup>. 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 infrastructure<sup>2</sup>. AWS Snowball Edge devices can run Amazon EC2 instances and AWS Lambda functions locally and sync data with AWS when network connectivity is available<sup>2</sup>.

#### QUESTION 17

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.

**Correct Answer: A, D, F**

**Section:**

**Explanation:**

Configuring AWS CloudTrail to log S3 data events will enable logging all activities for objects in the S3 bucket<sup>1</sup>. Data events are object-level API operations such as GetObject, DeleteObject, and PutObject<sup>1</sup>. 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 bucket<sup>2</sup>. EventBridge can route events from S3 to SNS, which can send emails to subscribers<sup>2</sup>. 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 way<sup>3</sup>. A lifecycle policy can transition the logs to a cheaper storage class such as Glacier or delete them after a specified period of time<sup>3</sup>.

#### QUESTION 18

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 able in a new Region Use the new table as a replica target tor 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 able m 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 m a new Region Use the new table as a replica target for DynamoDB global tables.

**Correct Answer: C**

**Section:**

**Explanation:**

[https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-global-table-stream-lambda/?nc1=h\\_ls](https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-global-table-stream-lambda/?nc1=h_ls)

#### QUESTION 19

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 DocumentD3 (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.

**Correct Answer: C**

**Section:**

**Explanation:**

On-demand capacity mode is the function of Dynamodb. <https://aws.amazon.com/blogs/news/running-spiky-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 20

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.

**Correct Answer: C**

**Section:**

**Explanation:**

Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Powered by Lustre, the world's most popular high-performance file system, FSx for Lustre offers shared storage with sub-ms latencies, up to terabytes per second of throughput, and millions of IOPS. FSx for Lustre file systems can also be linked to Amazon Simple Storage Service (S3) buckets, allowing you to access and process data concurrently from both a high-performance file system and from the S3 API.

#### QUESTION 21

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.

**Correct Answer: D**

**Section:**

**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 the Production OU. Then move the new account to the Production OU when AWS Config is done.

#### QUESTION 22

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.

**Correct Answer: C**

**Section:**

**Explanation:**

Aurora Auto Scaling enables your Aurora DB cluster to handle sudden increases in connectivity or workload. When the connectivity or workload decreases, Aurora Auto Scaling removes unnecessary Aurora Replicas so that you don't pay for unused provisioned DB instances

**QUESTION 23**

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 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
- B. 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.
- D. Configure subscription verification. Require the data customers to subscribe to the data product 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.

**Correct Answer: C**

**Section:**

**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 payment<sup>1</sup>.

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 stages<sup>2</sup>. 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 client<sup>3</sup>. 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 data<sup>4</sup>. 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.

<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 24**



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.

**Correct Answer: A**

**Section:**

**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 25

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.

**Correct Answer: A**

**Section:**

**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 block<sup>1</sup>. 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 connections<sup>2</sup>. 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 latency<sup>3</sup>. 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 cloud<sup>4</sup>. It does not provide any filtering capabilities.

<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 26

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.

**Correct Answer: B**

**Section:**

**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 27

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.

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

**Correct Answer: A**

**Section:**

**Explanation:**

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

#### QUESTION 28

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.

**Correct Answer: D**

**Section:**

**Explanation:**

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

#### QUESTION 29

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.

**Correct Answer: D**

**Section:**

**Explanation:**

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

#### QUESTION 30

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?
- B. 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.

- C. 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.
- D. 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.
- E. 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.

**Correct Answer: B**

**Section:**

**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 31

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.

**Correct Answer: A, D, F**

**Section:**

**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 automatically.

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 outage.

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 unreachable.

#### QUESTION 32

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.

**Correct Answer: C, E**

**Section:**

**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 33

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.

**Correct Answer: A, E**

**Section:**

**Explanation:**

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

### QUESTION 34

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.

**Correct Answer: A**

**Section:**

**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 35

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.

**Correct Answer: A**

**Section:**

**Explanation:**

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

#### QUESTION 36

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.

**Correct Answer: B**

**Section:**

**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 37**

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

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

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

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

**Correct Answer: C, E**

**Section:**

**Explanation:**

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

**QUESTION 38**

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 blog 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.

**Correct Answer: C**

**Section:**

**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 39**

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 interlace on each connection, and connect both private victual 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.

**Correct Answer: A**

**Section:**

**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 40

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

**Correct Answer: C**

**Section:**

**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 41

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 SLA. 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.

**Correct Answer: D**

**Section:**

**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 42

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.

**Correct Answer: B**

**Section:**

**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](https://docs.aws.amazon.com/secretsmanager/latest/userguide/integrating_cloudformation.html)

#### QUESTION 43

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

**Correct Answer: A, C**

**Section:**

**Explanation:**

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

#### QUESTION 44

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.

**Correct Answer: C, E, F**

**Section:**

**Explanation:**

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

#### QUESTION 45

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.

**Correct Answer: A**

**Section:**

**Explanation:**

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

#### QUESTION 46

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

**Correct Answer: A, C**

**Section:**

**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 47

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

**Correct Answer: C**

**Section:**

**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 AWSPrivateMarketplaceAdminFullAccess 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 procurement-manager-role and another SCP denies the permission to create an IAM role named procurement-manager-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 48

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": "AllowEC2",
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    },
    {
      "Sid": "AllowDynamoDB",
      "Effect": "Allow",
      "Action": "dynamodb:*",
      "Resource": "*"
    },
    {
      "Sid": "AllowS3",
      "Effect": "Allow",
      "Action": "s3:*",
      "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

**Correct Answer: B**

**Section:**

**Explanation:**

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_inheritance\\_auth.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_inheritance_auth.html)

#### QUESTION 49

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 multivalue answer routing policy with the IP addresses of all the EC2 instances. Recently, the app has been overwhelmed by large and sudden increases to 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.

**Correct Answer: D**

**Section:**

**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 50

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.

**Correct Answer: B**

**Section:**

**Explanation:**

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

#### QUESTION 51

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.

**Correct Answer: A**

**Section:**

**Explanation:**

According to the AWS documentation<sup>1</sup>, 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 52

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.

**Correct Answer: B**

**Section:**

**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 53

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.

**Correct Answer: B**

**Section:**

**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 54

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.

**Correct Answer: C**

**Section:**

**Explanation:**

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

#### QUESTION 55

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.

**Correct Answer: A, D**

**Section:**

**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 56

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.

**Correct Answer: A**

**Section:**

**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](https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf))

#### QUESTION 57

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.

**Correct Answer: B, D, F**

**Section:**

**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 itself<sup>1</sup>

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 connection<sup>23</sup>

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.

#### QUESTION 58

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.

**Correct Answer: A, D**

**Section:**

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

#### QUESTION 59

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

**Correct Answer: C, D, E**

**Section:**

**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 60

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

**Correct Answer: C**

**Section:**

**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 61

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

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

**Correct Answer: C**

**Section:**

**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 62

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.

**Correct Answer: D**

**Section:**

**Explanation:**

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

**QUESTION 63**

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.

**Correct Answer: A**

**Section:****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 64**

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.

**Correct Answer: A**

**Section:****Explanation:**

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

**QUESTION 65**

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.

**Correct Answer: C**

**Section:**

**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 66**

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.

**Correct Answer: A, B, E**

**Section:**

**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](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_tagging.html)

**QUESTION 67**

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 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 Quick Sight 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 0 DB instance to receive events Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards

**Correct Answer: A, D**

**Section:**

**Explanation:**

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

#### QUESTION 68

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

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

What should the solutions architect do to meet these requirements?

- A. Enable VPC Flow Logs. Use Amazon Athena to analyze the logs for traffic that can be removed. Ensure that security groups are Mocking traffic that is responsible for high costs.
- B. Add an interface VPC endpoint for Kinesis Data Streams to the 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.

**Correct Answer: D**

**Section:**

**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 69

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.

**Correct Answer: D**

**Section:**

**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 70

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.

**Correct Answer: A, D, E**

**Section:**

**Explanation:**

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

#### QUESTION 71

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.

**Correct Answer: A, C, D**

**Section:**

**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 72

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.

**Correct Answer: D**

**Section:**

**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 73

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

**Correct Answer: A, E**

**Section:**

**Explanation:**

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

#### QUESTION 74

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 SecretsManagerReadOnly 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 SecretsManagerReadOnly 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 SecretsManagerReadOnly 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

SecretsManagerReadWnte 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 SecretsManagerReadWnte 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 SecretsManagerReadWnte policy that is attached to the IAM role

**Correct Answer: B**

**Section:**

**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 SecretsManagerReadWnte 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](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_simulator.html)

#### QUESTION 75

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. 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.

**Correct Answer: C**

**Section:**

**Explanation:**

Generally, unstructured data should be converted structured data before querying them. AWS Glue cando that. <https://docs.aws.amazon.com/glue/latest/dg/schema-relationalize.html><https://docs.aws.amazon.com/athena/latest/ug/glue-athena.html>

#### QUESTION 76

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.

**Correct Answer: B**

**Section:****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.

**QUESTION 77**

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.

**Correct Answer: D**

**Section:****Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/vpc-analyze-inbound-traffic-nat-gateway/> by Cloudxie says 'select appropriate log'

**QUESTION 78**

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.

**Correct Answer: C**

**Section:****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 79**

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.

**Correct Answer: B**

**Section:**

**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 requests<sup>1</sup>. 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 session<sup>2</sup>.

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 instances<sup>3</sup>. 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 requests<sup>4</sup>. However, this algorithm does not take into account customer sessions or weights.

<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 80

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.

**Correct Answer: D**

**Section:**

**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.

<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 81

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. 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.

**Correct Answer: D**

**Section:**

**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<sup>1</sup>. 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<sup>2</sup>. 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<sup>3</sup>.

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.

<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 82

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.

**Correct Answer: A**

**Section:**

**Explanation:**

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html>

### QUESTION 83

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.

**Correct Answer: A**

**Section:**

**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 files<sup>1</sup>. 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 durability<sup>2</sup>. 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 balance<sup>3</sup>. 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.

a. 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.

<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 84

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.

**Correct Answer: C**

**Section:**

**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 blocked<sup>1</sup>. 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 patterns<sup>2</sup>. 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 database<sup>3</sup>.

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 rules<sup>4</sup>.

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.

<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 85

A company is updating an application that customers use to make online orders. The number of attacks on the application by bad actors has increased recently.

The company will host the updated application on an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use Amazon DynamoDB to store application data. A public Application Load Balancer (ALB) will provide end users with access to the application. The company must prevent prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack.

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



- A. Create an Amazon CloudFront distribution with the ALB as the origin. Add a custom header and random value 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.

**Correct Answer: A, E**

**Section:**

**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 environment<sup>1</sup>. 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 URLs<sup>2</sup>. 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 blocked<sup>3</sup>. 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 addresses<sup>4</sup>. 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.

<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 86

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.

**Correct Answer: A**

**Section:**

**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-month<sup>1</sup>. 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 days<sup>2</sup>, otherwise they incur a pro-rated charge equal to the storage charge for the remaining days<sup>1</sup>. 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 87

A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. 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.

**Correct Answer: C**

**Section:**

**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 sources<sup>1</sup>. 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 backend<sup>1</sup>. AWS AppSync uses WebSockets to establish and maintain a secure connection between the clients and the API endpoint<sup>2</sup>. Therefore, using AWS AppSync and leveraging WebSockets is a suitable design to reduce comment latency and improve user experience.

#### QUESTION 88

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.

**Correct Answer: B, D**

**Section:**

**Explanation:**

AWS Trusted Advisor is a service that provides real-time guidance to help users provision their resources following AWS best practices<sup>1</sup>. 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 metrics<sup>1</sup>. 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 workloads<sup>2</sup>. Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWS Fargate, and Lambda functions<sup>2</sup>. 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 instances<sup>3</sup>. It then recommends optimal instance types based on price-performance trade-offs.

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 plan<sup>1</sup>.

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 enabled<sup>3</sup>.

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 term<sup>4</sup>. Savings Plans do not affect the configuration or utilization of EC2 instances.

#### QUESTION 89

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.

**Correct Answer: C**

**Section:**

**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 deletion<sup>1</sup>. 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 resources<sup>2</sup>. 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 Lambda<sup>2</sup>.

S3 Lifecycle is a feature that allows users to manage their objects so that they are stored cost-effectively throughout their lifecycle<sup>3</sup>. 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 needed<sup>4</sup>. 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 90

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.

**Correct Answer: B**

**Section:**

**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](https://docs.aws.amazon.com/fr_fr/kms/latest/developerguide/services-s3.html)

#### QUESTION 91

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.

**Correct Answer: B, D**

**Section:**

#### QUESTION 92

A company's solutions architect is analyzing costs of a multi-application environment. The environment is deployed across multiple Availability Zones in a single AWS Region. After a recent acquisition, the company manages two organizations in AWS Organizations. The company has created multiple service provider applications as AWS PrivateLink-powered VPC endpoint services in one organization. The company has created multiple service consumer applications in the other organization.

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

Which guidelines meet these requirements? (Select TWO.)

- A. Use AWS Resource Access Manager to share the subnets that host the service provider applications with other accounts in the organization.
- B. Place the service provider applications and the service consumer applications in AWS accounts in the same organization.
- C. Turn off cross-zone load balancing for the Network Load Balancer in all service provider application deployments.
- D. Ensure that service consumer compute resources use the Availability Zone-specific endpoint service by using the endpoint's local DNS name.
- E. Create a Savings Plan that provides adequate coverage for the organization's planned inter-Availability Zone data transfer usage.

**Correct Answer: C, D**

**Section:**

**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 inter-Availability 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 93

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.

**Correct Answer: A**

**Section:**

**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 94

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.

**Correct Answer: A, C**

**Section:**

**Explanation:**



<https://joe.blog.freemansoft.com/2020/04/protect-data-in-cloud-with-s3-access.html>

#### QUESTION 95

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.

**Correct Answer: B, C, E**

**Section:**

**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_introduction.html)

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scp.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/IAM/latest/UserGuide/id_roles.html)

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

#### QUESTION 96

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).

**Correct Answer: A**

**Section:**

**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 97

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 fallback, 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.

**Correct Answer: D**

**Section:**

#### QUESTION 98

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.

**Correct Answer: B, C, E**

**Section:**

**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 99**

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.

**Correct Answer: B, E, F**

**Section:**

**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](https://docs.aws.amazon.com/organizations/latest/APIReference/API_RemoveAccountFromOrganization.html)

**QUESTION 100**

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) file 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.

**Correct Answer: D**

**Section:**

**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 101**

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 endpoint 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.

**Correct Answer: A**

**Section:**

**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 102

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

**Correct Answer: A**

**Section:**

**Explanation:**

<https://aws.amazon.com/blogs/mt/implement-aws-resource-tagging-strategy-using-aws-tag-policies-and-service-control-policies-scps/>

#### QUESTION 103

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.

**Correct Answer: B**

**Section:**

**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.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.SQLServer.html)

[https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Target.S3.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html)

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_SQLServer.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServer.html)

#### QUESTION 104

A company has set up its entire infrastructure on AWS. The company uses Amazon EC2 instances to host its ecommerce website and uses Amazon S3 to store static data. Three engineers at the company handle the cloud administration and development through one AWS account. Occasionally, an engineer alters an EC2 security group configuration of another engineer and causes noncompliance issues in the environment.

A solutions architect must set up a system that tracks changes that the engineers make. The system must send alerts when the engineers make noncompliant changes to the security settings for the EC2 instances.

What is the FASTEST way for the solutions architect to meet these requirements?

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

**Correct Answer: D**

**Section:**

**Explanation:**

<https://aws.amazon.com/es/blogs/industries/how-to-monitor-alert-and-remediate-non-compliant-hipaa-findings-on-aws/>

#### QUESTION 105

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.

**Correct Answer: A, B**

**Section:**

**Explanation:**

<https://medium.com/@it.melnichenko/invoke-a-lambda-across-multiple-aws-accounts-8c094b2e70be>

**QUESTION 106**

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. 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.

**Correct Answer: A**

**Section:**

**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.

<https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/>

<https://aws.amazon.com/autoscaling/>

<https://aws.amazon.com/route53/>

**QUESTION 107**

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.

**Correct Answer: A**

**Section:**

**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 Organizations<sup>1</sup>. Creating a backup plan that specifies the frequency and retention requirements will enable taking snapshots every 6 hours and retaining them for 30 days<sup>2</sup>. Adding a tag to the DB instances will enable applying the backup plan by using tags<sup>2</sup>.



Using AWS Backup to monitor the status of the backups will enable having a consolidated view of the health of the RDS snapshots<sup>1</sup>.

#### QUESTION 108

A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes.

The company must maintain backups in a separate AWS Region. The company must be able to recover the EC2 instances and their configuration within 1 business day, with loss of no more than 1 day's worth of data. The company has limited staff and needs a backup solution that optimizes operational efficiency and cost. The company already has created an AWS CloudFormation template that can deploy the required network configuration in a secondary Region.

Which solution will meet these requirements?

- A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbooks. Copy the snapshots to the secondary Region. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.
- B. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volumes. In the event of a failure, launch the CloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.
- C. Use AWS Backup to create a scheduled daily backup plan for the EC2 instances. Configure the backup task to copy the backups to a vault in the secondary Region. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.
- D.

**Correct Answer: C**

**Section:**

**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 volumes<sup>1</sup>. Configuring the backup task to copy the backups to a vault in the secondary Region will enable maintaining backups in a separate Region<sup>1</sup>. In the event of a failure, launching the CloudFormation template will enable deploying the network configuration in the secondary Region<sup>2</sup>. Restoring the instance volumes and configurations from the backup vault will enable recovering the EC2 instances and their data<sup>1</sup>. Transferring usage to the secondary Region will enable resuming operations<sup>2</sup>.

#### QUESTION 109

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. 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.

**Correct Answer: C**

**Section:**

**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 names<sup>1</sup>. 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 have<sup>2</sup>. 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 resources<sup>3</sup>. 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)<sup>4</sup>.

Permissions boundaries cannot be used to specify fine-grained access control for DynamoDB items and attributes.

<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/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/id_roles.html)

[https://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies\\_boundaries.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html)

**QUESTION 110**

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.

**Correct Answer: C**

**Section:**

**Explanation:**

Adding an accelerator in AWS Global Accelerator will enable improving the performance of the APIs for local and global users<sup>1</sup>. 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 policies<sup>1</sup>. Configuring the ALB as the origin will enable connecting the accelerator to the ALB that exposes the APIs<sup>2</sup>. AWS Global Accelerator supports non-standard REST methods such as LINK, UNLINK, LOCK, and UNLOCK<sup>3</sup>.

**QUESTION 111**

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.

**Correct Answer: D**

**Section:**

**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 112

A company manufactures smart vehicles. The company uses a custom application to collect vehicle data. The vehicles use the MQTT protocol to connect to the application.

The company processes the data in 5-minute intervals. The company then copies vehicle telematics data to on-premises storage. Custom applications analyze this data to detect anomalies.

The number of vehicles that send data grows constantly. Newer vehicles generate high volumes of data. The on-premises storage solution is not able to scale for peak traffic, which results in data loss. The company must modernize the solution and migrate the solution to AWS to resolve the scaling challenges.

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

- A. Use AWS IoT Greengrass to send the vehicle data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create an Apache Kafka application to store the data in Amazon S3. Use a pretrained model in Amazon SageMaker to detect anomalies.
- B. Use AWS IoT Core to receive the vehicle data. Configure rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3. Create an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies.
- C. Use AWS IoT FleetWise to collect the vehicle data. Send the data to an Amazon Kinesis data stream. Use an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use the built-in machine learning transforms in AWS Glue to detect anomalies.
- D. Use Amazon MQ for RabbitMQ to collect the vehicle data. Send the data to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Lookout for Metrics to detect anomalies.

**Correct Answer: B**

**Section:**

**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 113

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.

**Correct Answer: A, C, E**

**Section:**

**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)<sup>1</sup>. 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 bucket<sup>2</sup>. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS<sup>3</sup>.

#### QUESTION 114

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.

**Correct Answer: C, D, E**

**Section:**

**Explanation:**

Enabling all features for the organization will enable using AWS Firewall Manager to centrally configure and manage firewall rules across multiple AWS accounts<sup>1</sup>. 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 vulnerabilities<sup>2</sup>. AWS Firewall Manager supports AWS WAF rules that can help protect against common web exploits such as SQL injection and cross-site scripting<sup>3</sup>. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS.

#### QUESTION 115

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.



**Correct Answer: C**

**Section:**

**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 servers<sup>1</sup>.
- \* Use AWS AppSync to create a scalable GraphQL API that connects to your database and other data sources<sup>1</sup>.
- \* Use Amazon SQS to decouple and scale your order processing microservices<sup>1</sup>.
- \* Use AWS Lambda to run code for your business logic without provisioning or managing servers<sup>1</sup>.
- \* Use an Amazon SQS dead-letter queue to retain messages that can't be processed by your Lambda function<sup>1</sup>.

#### QUESTION 116

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.

**Correct Answer: B**

**Section:**

**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 minutes<sup>1</sup>. 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 minutes<sup>2</sup>. 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 Region<sup>3</sup>. By using infrastructure as code (IaC), the company can provision the new infrastructure in the DR Region in an automated and consistent way<sup>4</sup>. 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.

a. It is useful for sharing query results and views with other users, but not for replicating or recovering data.

<https://aws.amazon.com/disaster-recovery/>

<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://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html#USER_ReadRepl.XRgn)

<https://aws.amazon.com/cloudformation/>

<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-exchange/>

<https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html>

#### QUESTION 117

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

**Correct Answer: C**

**Section:**

**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 Region<sup>1</sup>. 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 latency<sup>2</sup>. 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 warehouse<sup>3</sup>. 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.

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS\\_Fea\\_Regions\\_DB-eng.Feature.CrossRegionReadReplicas.html](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 118

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.

**Correct Answer: A**

**Section:**

**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 discount<sup>1</sup>. 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.

a. EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings<sup>2</sup>. 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.

<https://aws.amazon.com/ec2/pricing/reserved-instances/>

<https://aws.amazon.com/savingsplans/>

#### QUESTION 119

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 Cloud Formation 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 Cloud Formation 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.

**Correct Answer: B, D, E**

**Section:**

**Explanation:**

#### QUESTION 120

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 cross-account 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.

**Correct Answer: D**

**Section:**

**Explanation:**

Cross account role should be created in destination(member) account. The role has trust entity to master account.

#### QUESTION 121

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.

**Correct Answer: D**

**Section:**

**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 122

A company maintains a restaurant review website. The website is a single-page application where files are stored in Amazon S3 and delivered using Amazon CloudFront. The company receives several fake postings every day that are manually removed.

The security team has identified that most of the fake posts are from bots with IP addresses that have a bad reputation within the same global region. The team needs to create a solution to help restrict the bots from accessing the website. Which strategy should a solutions architect use?

- A. Use AWS Firewall Manager to control the CloudFront distribution security settings. Create a geographical block rule and associate it with Firewall Manager.
- B. Associate an AWS WAF web ACL with the CloudFront distribution. Select the managed Amazon IP reputation rule group for the web ACL with a deny action.
- C. Use AWS Firewall Manager to control the CloudFront distribution security settings. Select the managed Amazon IP reputation rule group and associate it with Firewall Manager with a deny action.
- D. Associate an AWS WAF web ACL with the CloudFront distribution. Create a rule group for the web ACL with a geographical match statement with a deny action.

**Correct Answer: B**

**Section:**

**Explanation:**

IP reputation rule groups allow you to block requests based on their source. Choose one or more of these rule groups if you want to reduce your exposure to BOTS!!!! traffic or exploitation attempts The Amazon IP reputation list rule group contains rules that are based on Amazon internal threat intelligence. This is useful if you would like to block IP addresses typically associated with bots or other threats. Inspects for a list of IP addresses that have been identified as bots by Amazon threat intelligence.

#### QUESTION 123

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.

**Correct Answer: C**

**Section:**

**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 124

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.

**Correct Answer: B**

**Section:**

**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 S3<sup>1</sup>. AWS DMS supports homogeneous migrations such as migrating from one Aurora MySQL DB cluster to another, as well as heterogeneous migrations between different database platforms<sup>2</sup>. AWS DMS also supports cross-account migrations, as long as the source and target databases are in the same AWS Region<sup>3</sup>.

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 inconsistency<sup>4</sup>.

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.

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 125

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 OU.
- 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.

**Correct Answer: A**

**Section:**

**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.

#### QUESTION 126

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 EFS file 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 server to place files into an Amazon FSx for Lustre 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 Lustre endpoint instead.

**Correct Answer: A**

**Section:**

**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.

#### QUESTION 127

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.

**Correct Answer: D**

**Section:****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 128**

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.

**Correct Answer: A**

**Section:****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 129**

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.

**Correct Answer: A**

**Section:****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

[www.VCEplus.io](http://www.VCEplus.io)