Number: SAP-C02
Passing Score: 800
Time Limit: 120
File Version: 31.0

Exam Code: SAP-C02

Exam Name: AWS Certified Solutions Architect - Professional

Exam A

QUESTION 1
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 Cloud Formation 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 2
A retail company needs to provide a series of data files to another company, which is its business partner These files are saved in an Amazon S3 bucket under Account
A which belongs to the retail company. The business partner company wants one of its 1AM users. User_DataProcessor. to access the files from its own AWS account (Account B).
Which combination of steps must the companies take so that User_DataProcessor can access the S3 bucket successfully? (Select TWO.)

A. Turn on the cross-origin resource sharing (CORS) feature for the S3 bucket in Account
B. In Account A. set the S3 bucket policy to the following:

```
{
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
    },
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::AccountABucketName/*"
    ]
}
```

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

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

```
{
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
    },
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::AccountABucketName/*"
    ]
}
```

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

Correct Answer: C, D
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-s3/

QUESTION 3
A utility company wants to collect usage data every 5 minutes from its smart meters to facilitate time-of-use metering. When a meter sends data to AWS. the data is sent to Amazon API Gateway, processed by an AWS Lambda function, and stored in an Amazon DynamoDB table. During the pilot phase, the Lambda functions took from 3 to 5 seconds to complete.
As more smart meters are deployed, the engineers notice the Lambda functions are taking from 1 to 2 minutes to complete. The functions are also increasing in duration as new types of metrics are collected from the devices.
There are many ProvisionedThroughputExceededException errors while performing PUT operations on DynamoDB. and there are also many TooManyRequestsException errors from Lambda.
Which combination of changes will resolve these issues? (Select TWO)

A. Increase the write capacity units to the DynamoDB table.
B. Increase the memory available to the Lambda functions
C. Increase the payload size from the smart meters to send more data.
D. Stream the data into an Amazon Kinesis data stream from API Gateway and process the data in batches.
E. Collect data in an Amazon SOS FIFO queue, which triggers a Lambda function to process each message.

Correct Answer: A, B
Section:

QUESTION 4
A company is running a serverless ecommerce application on AWS. The application uses Amazon API Gateway to invoke AWS Lambda Java functions. The Lambda functions connect to an Amazon RDS for MySQL database to store data.
During a recent sale event, a sudden increase in web traffic resulted in poor API performance and database connection failures. The company needs to implement a solution to minimize the latency for the Lambda functions and to support bursts in traffic.
Which solution will meet these requirements with the LEAST amount of change to the application?

A. Update the code of the Lambda functions so that the Lambda functions open the database connection outside of the function handler. Increase the provisioned concurrency for the Lambda functions.
B. Create an RDS Proxy endpoint for the database. Store database secrets in AWS Secrets Manager. Set up the required 1AM permissions. Update the Lambda functions to connect to the RDS Proxy endpoint. Increase the provisioned concurrency for the Lambda functions.

C. Create a custom parameter group. Increase the value of the max_connections parameter. Associate the custom parameter group with the RDS DB instance and schedule a reboot. Increase the reserved concurrency for the Lambda functions.
D. Create an RDS Proxy endpoint for the database. Store database secrets in AWS Secrets Manager. Set up the required 1AM permissions. Update the Lambda functions to connect to the RDS Proxy endpoint. Increase the reserved concurrency for the Lambda functions.

Correct Answer: B
Section:

QUESTION 5
A solutions architect is redesigning a three-tier application that a company hosts on premises. The application provides personalized recommendations based on user profiles. The company already has an AWS account and has configured a VPC to host the application.
The frontend is a Java-based application that runs in on-premises VMs. The company hosts a personalization model on a physical application server and uses TensorFlow to implement the model. The personalization model uses artificial intelligence and machine learning (AI/ML). The company stores user information in a Microsoft SQL Server database. The web application calls the personalization model, which reads the user profiles from the database and provides recommendations.
The company wants to migrate the redesigned application to AWS.
Which solution will meet this requirement with the LEAST operational overhead?

A. Use AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
B. Export the personalization model. Store the model artifacts in Amazon S3. Deploy the model to Amazon SageMaker and create an endpoint. Host the Java application in AWS Elastic Beanstalk. Use AWS Database Migration Service {AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
C. Use AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling groups. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to an EC2 instance.
D. Containerize the personalization model and the Java application. Use Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy the model and the application to Amazon EKS Host the node groups in a VPC. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.

Correct Answer: B
Section:
Explanation:
Amazon SageMaker is a fully managed machine learning service that allows users to build, train, and deploy machine learning models quickly and easily1.Users can export their existing TensorFlow models and store the model artifacts in Amazon S3, a highly scalable and durable object storage service2.Users can then deploy the model to Amazon SageMaker and create an endpoint that can be invoked by the web application to provide recommendations3. This way, the solution can leverage the AI/ML capabilities of Amazon SageMaker without having to rewrite the personalization model.
AWS Elastic Beanstalk is a service that allows users to deploy and manage web applications without worrying about the infrastructure that runs those applications. Users can host their Java application in AWS Elastic Beanstalk and configure it to communicate with the Amazon SageMaker endpoint. This way, the solution can reduce the operational overhead of managing servers, load balancers, scaling, and application health monitoring.
AWS Database Migration Service (AWS DMS) is a service that helps users migrate databases to AWS quickly and securely. Users can use AWS DMS to migrate their SQL Server database to Amazon RDS for SQL Server, a fully managed relational database service that offers high availability, scalability, security, and compatibility. This way, the solution can reduce the operational overhead of managing database servers, backups, patches, and upgrades.
Option A is incorrect because using AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS is not cost-effective or scalable. AWS SMS is a service that helps users migrate on-premises workloads to AWS. However, for this use case, migrating the physical application server and the web application VMs to AWS will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.
Option C is incorrect because using AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling groups is not cost-effective or scalable. AWS Application Migration Service is a service that helps users migrate applications from on-premises or other clouds to AWS without making any changes to their applications. However, for this use case, migrating the personalization model and VMs to EC2 instances will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.
Option D is incorrect because containerizing the personalization model and the Java application and using Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy them to Amazon EKS is not necessary or cost-effective. Amazon EKS is a service that allows users to run Kubernetes on AWS without needing to install, operate, and maintain their own Kubernetes control plane or nodes. However, for this use case, containerizing and deploying the personalization model and the Java application will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk. Moreover, using S3 Glacier Deep Archive as a storage class for images will incur a high retrieval fee and latency for accessing them.

QUESTION 6

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

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

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-traffic-shifting/

QUESTION 7
A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.
How can the company prevent users from accidentally deleting data in this way?

A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
C. Modify 1AM policies to deny deleting RDS and EBS resources that are tagged with an 'awsrcloudformation: stack-name' tag.
D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Correct Answer: A
Section:
Explanation:
With the DeletionPolicy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default. To keep a resource when its stack is deleted, specify Retain for that resource. You can use retain for any resource. For example, you can retain a nested stack, Amazon S3 bucket, or EC2 instance so that you can continue to use or modify those resources after you delete their stacks. https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html

QUESTION 8
A company runs an intranet application on premises. The company wants to configure a cloud backup of the application. The company has selected AWS Elastic Disaster Recovery for this solution.
The company requires that replication traffic does not travel through the public internet. The application also must not be accessible from the internet. The company does not want this solution to consume all available network bandwidth because other applications require bandwidth.
Which combination of steps will meet these requirements? (Select THREE.)

A. Create a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway.
B. Create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway.
C. Create an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network.
D. Create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network.
E. During configuration of the replication servers, select the option to use private IP addresses for data replication.
F. During configuration of the launch settings for the target servers, select the option to ensure that the Recovery instance's private IP address matches the source server's private IP address.

Correct Answer: B, D, E
Section:

Explanation:
AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery1. Users can set up AWS DRS on their source servers to initiate secure data replication to a staging area subnet in their AWS account, in the AWS Region they select. Users can then launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time.
To configure a cloud backup of the application with AWS DRS, users need to create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway.A VPC is a logically isolated section of the AWS Cloud where users can launch AWS resources in a virtual network that they define2.A public subnet is a subnet that has a route to an internet gateway3.A virtual private gateway is the VPN concentrator on the Amazon side of the Site-to-Site VPN connection4. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in the VPC and the internet. Users need to create at least two public subnets for redundancy and high availability. Users need to create a virtual private gateway and attach it to the VPC to enable VPN connectivity between the on-premises network and the target AWS network. Users need to create an internet gateway and attach it to the VPC to enable internet access for the replication servers.
To ensure that replication traffic does not travel through the public internet, users need to create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network. AWS Direct Connect is a service that establishes a dedicated network connection from an on-premises network to one or more VPCs. A Direct Connect gateway is a globally available resource that allows users to connect multiple VPCs across different Regions to their on-premises networks using one or more Direct Connect connections. Users need to create an AWS Direct Connect connection between their on-premises network and an AWS Region. Users need to create a Direct Connect gateway and associate it with their VPC and their Direct Connect connection.
To ensure that the application is not accessible from the internet, users need to select the option to use private IP addresses for data replication during configuration of the replication servers. This option configures the replication servers with private IP addresses only, without assigning any public IP addresses or Elastic IP addresses. This way, the replication servers can only communicate with other resources within the VPC or through VPN connections.
Option A is incorrect because creating a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway is not necessary or cost-effective.A private subnet is a subnet that does not have a route to an internet gateway3. A NAT gateway is a highly available, managed Network Address Translation (NAT) service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances. Users do not need to create private subnets or NAT gateways for this use case, as they can use public subnets with private IP addresses for data replication.
Option C is incorrect because creating an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network will not ensure that replication traffic does not travel through the public internet.A Site-to-Site VPN connection consists of two VPN tunnels between an on-premises customer gateway device and a virtual private gateway in your VPC4. The VPN tunnels are encrypted using IPSec protocols, but they still use public IP addresses for communication. Users need to use AWS Direct Connect instead of Site-to-Site VPN for this use case.
Option F is incorrect because selecting the option to ensure that the Recovery instance's private IP address matches the source server's private IP address during configuration of the launch settings for the target servers will not ensure that the application is not accessible from the internet. This option configures the Recovery instance with an identical private IP address as its source server when launched in drills or recovery mode. However, this option does not prevent assigning public IP addresses or Elastic IP addresses to the Recovery instance. Users need to select the option to use private IP addresses for data replication instead.

QUESTION 9
A software company has deployed an application that consumes a REST API by using Amazon API Gateway. AWS Lambda functions, and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.
A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.
What should the solutions architect recommend to improve the customer experience?

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

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/aws-batch-requests-error/
https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-429-limit/

QUESTION 10
A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.
A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.

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

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

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 11
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 12
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 13
A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application's data. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record.
The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an active-active strategy.
What should a solutions architect recommend to meet these requirements?

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

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

C. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Region. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs. Remove the read replica. Replace the read replica with a standalone RDS DB instance. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.

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

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 14
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 automatically1
Option D is correct because modifying the DB instance to create a Multi-AZ deployment that extends across two Availability Zones can improve the availability and durability of the RDS for MariaDB database.Multi-AZ deployments provide enhanced data protection and minimize downtime by automatically failing over to a standby replica in another Availability Zone in case of a planned or unplanned outage4
Option F is correct because creating a replication group for the ElastiCache for Redis cluster and enabling Multi-AZ on the cluster can improve the availability and fault tolerance of the in-memory data store. A replication group consists of a primary node and up to five read-only replica nodes that are synchronized with the primary node using asynchronous replication.Multi-AZ allows automatic failover to one of the replicas if the primary node fails or becomes unreachable6

QUESTION 15
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 16
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 17
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 VPUpdate 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 18
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 OuickSight 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 19
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 20
A company is using multiple AWS accounts The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A The company's applications and databases are running in Account B.
A solutions architect win deploy a two-net application In a new VPC To simplify the configuration, the db.example com CNAME record set tor the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.
During deployment, the application failed to start. Troubleshooting revealed that db.example com is not resolvable on the Amazon EC2 instance The solutions architect confirmed that the record set was created correctly in Route 53.
Which combination of steps should the solutions architect take to resolve this issue? (Select TWO)

A. Deploy the database on a separate EC2 instance in the new VPC Create a record set for the instance's private IP in the private hosted zone
B. Use SSH to connect to the application tier EC2 instance Add an RDS endpoint IP address to the /eto/resolv.conf file
C. Create an authorization lo associate the private hosted zone in Account A with the new VPC In Account B
D. Create a private hosted zone for the example.com domain m Account B Configure Route 53 replication between AWS accounts
E. Associate a new VPC in Account B with a hosted zone in Account 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 21
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 ScaSng group The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature 'or Moggers 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 win resolve the issues for users?

A. Reconfigure Amazon EFS to enable maximum I/O.

B. Update the Nog site to use instance store volumes tor storage. Copy the site contents to the volumes at launch and to Amazon S3 al 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 22
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 23
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 24
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 25
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

QUESTION 26
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 27
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 28
A company that has multiple AWS accounts is using AWS Organizations. The company's AWS accounts host VPCs, Amazon EC2 instances, and containers.
The company's compliance team has deployed a security tool in each VPC where the company has deployments. The security tools run on EC2 instances and send information to the AWS account that is dedicated for the

compliance team. The company has tagged all the compliance-related resources with a key of ''costCenter'' and a value or ''compliance''.
The company wants to identify the cost of the security tools that are running on the EC2 instances so that the company can charge the compliance team's AWS account. The cost calculation must be as accurate as possible.
What should a solutions architect do to meet these requirements?

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

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 29
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. Prom 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 automatical^/ 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 automatical^ 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 30
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 31
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

QUESTION 32
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 33
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 34
A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.
The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS.
Which data migration strategy should the company use?

A. Use the file gateway option in AWS Storage Gateway to replace the existing Windows file server, and point the existing file share to the new file gateway.
B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.

C. Use AWS Data Pipeline to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).
D. Use AWS DataSync to schedule a daily task lo replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS),

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/storagegateway/file/
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html
https://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os-windows-server

QUESTION 35
A company's solutions architect is reviewing a web application that runs on AWS. The application references static assets in an Amazon S3 bucket in the us-east-1 Region. The company needs resiliency across multiple AWS Regions. The company already has created an S3 bucket in a second Region.
Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html

QUESTION 36
A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database A solutions architect must design a scalable and highly available solution to meet the demand of 200000 daily users.
Which steps should the solutions architect take to design an appropriate solution?

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

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

QUESTION 37
A company is using AWS Organizations lo 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://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-accounts-and-regions/

QUESTION 38
A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises intra structure that consists of physical machines and VMs that host numerous applications.
The company must capture details about the system configuration. system performance. running processure and network coi.net lions of its o. -premises ,on boards. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.
Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

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

Correct Answer: A, D, E
Section:
Explanation:
https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html https://docs.aws.amazon.com/migrationhub/latest/ug/ec2-recommendations.html

QUESTION 39
A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.
The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 ' of data from an S3 bucket each day.
The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations.
Which solution will meet these requirements?

A. Replace the NAT gateways with NAT instances. In the VPC route table, create a route from the private subnets to the NAT instances.
B. Move the EC2 instances to the public subnets. Remove the NAT gateways.
C. Set up an S3 gateway VPC endpoint in the VPC. Attach an endpoint policy to the endpoint to allow the required actions on the S3 bucket.
D. Attach an Amazon Elastic File System (Amazon EFS) volume to the EC2 instances. Host the image on the EFS volume.

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

QUESTION 40
A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load tor the rest of the week. The access pattern includes many more writes to the table than reads of the table.
A solutions architect needs to implement a solution to minimize the cost of the table.
Which solution will meet these requirements?

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

Correct Answer: A
Section:
Explanation:

QUESTION 41
A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with the number of files rapidly declining after business hours.
What is the MOST cost-effective migration recommendation?

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

Correct Answer: D
Section:
Explanation:
https://aws.amazon.com/blogs/compute/operating-lambda-performance-optimization-part-1/

QUESTION 42
A company is using Amazon OpenSearch Service to analyze data. The company loads data into an OpenSearch Service cluster with 10 data nodes from an Amazon S3 bucket that uses S3 Standard storage. The data resides in the cluster for 1month for read-only analysis. After 1month, the company deletes the index that contains the data from the cluster. For compliance purposes, the company must retain a copy of all input data.
The company is concerned about ongoing costs and asks a solutions architect to recommend a new solution.
Which solution will meet these requirements MOST cost-effectively?

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

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

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

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

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

QUESTION 43
A company has 10 accounts that are part of an organization in AWS Organizations AWS Config is configured in each account All accounts belong to either the Prod OU or the NonProd OU
The company has set up an Amazon EventBridge rule in each AWS account to notify an Amazon Simple Notification Service (Amazon SNS) topic when an Amazon EC2 security group inbound rule is created with 0.0.0.0/0 as the source The company's security team is subscribed to the SNS topic
For all accounts in the NonProd OU the security team needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source
Which solution will meet this requirement with the LEAST operational overhead?

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

Correct Answer: D
Section:
Explanation:
This solution will meet the requirement with the least operational overhead because it directly denies the creation of the security group inbound rule with 0.0.0.0/0 as the source, which is the exact requirement. Additionally, it does not require any additional steps or resources such as invoking a Lambda function or adding a Config rule.
An SCP (Service Control Policy) is a policy that you can use to set fine-grained permissions for your AWS accounts within your organization. You can use SCPs to set permissions for the root user of an account and to delegate permissions to IAM users and roles in the accounts. You can use SCPs to set permissions that allow or deny access to specific services, actions, and resources.
To implement this solution, you would need to create an SCP that denies the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0. This SCP would then be applied to the NonProd OU. This would ensure that any security group inbound rule that includes 0.0.0.0/0 as the source will be denied, thus meeting the requirement.
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_condition-keys.html

QUESTION 44
A company hosts a Git repository in an on-premises data center. The company uses webhooks to invoke functionality that runs in the AWS Cloud. The company hosts the webhook logic on a set of Amazon EC2 instances in an Auto Scaling group that the company set as a target for an Application Load Balancer (ALB). The Git server calls the ALB for the configured webhooks. The company wants to move the solution to a serverless architecture.
Which solution will meet these requirements with the LEAST operational overhead?

A. For each webhook, create and configure an AWS Lambda function URL. Update the Git servers to call the individual Lambda function URLs.
B. Create an Amazon API Gateway HTTP API. Implement each webhook logic in a separate AWS Lambda function. Update the Git servers to call the API Gateway endpoint.
C. Deploy the webhook logic to AWS App Runner. Create an ALB, and set App Runner as the target. Update the Git servers to call the ALB endpoint.

D. Containerize the webhook logic. Create an Amazon Elastic Container Service (Amazon ECS) cluster, and run the webhook logic in AWS Fargate. Create an Amazon API Gateway REST API, and set Fargate as the target. Update the Git servers to call the API Gateway endpoint.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/solutions/implementations/git-to-s3-using-webhooks/
https://medium.com/mindorks/building-webhook-is-easy-using-aws-lambda-and-api-gateway-56f5e5c3a596

QUESTION 45
A company is planning to migrate 1,000 on-premises servers to AWS. The servers run on several VMware clusters in the company's data center. As part of the migration plan, the company wants to gather server metrics such as CPU details, RAM usage, operating system information, and running processes. The company then wants to query and analyze the data.
Which solution will meet these requirements?

A. Deploy and configure the AWS Agentless Discovery Connector virtual appliance on the on-premises hosts. Configure Data Exploration in AWS Migration Hub. Use AWS Glue to perform an ETL job against the data. Query the data by using Amazon S3 Select.
B. Export only the VM performance information from the on-premises hosts. Directly import the required data into AWS Migration Hub. Update any missing information in Migration Hub. Query the data by using Amazon QuickSight.
C. Create a script to automatically gather the server information from the on-premises hosts. Use the AWS CLI to run the put-resource-attributes command to store the detailed server data in AWS Migration Hub. Query the data directly in the Migration Hub console.
D. Deploy the AWS Application Discovery Agent to each on-premises server. Configure Data Exploration in AWS Migration Hub. Use Amazon Athena to run predefined queries against the data in Amazon S3.

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

QUESTION 46
A company is building a serverless application that runs on an AWS Lambda function that is attached to a VPC. The company needs to integrate the application with a new service from an external provider. The external provider supports only requests that come from public IPv4 addresses that are in an allow list.
The company must provide a single public IP address to the external provider before the application can start using the new service.
Which solution will give the application the ability to access the new service?

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

Correct Answer: A
Section:
Explanation:
This solution will give the Lambda function access to the internet by routing its outbound traffic through the NAT gateway, which has a public Elastic IP address. This will allow the external provider to whitelist the single public IP address associated with the NAT gateway, and enable the application to access the new service
Deploying a NAT gateway and associating an Elastic IP address with it, and then configuring the VPC to use the NAT gateway, will give the application the ability to access the new service. This is because the NAT gateway will be the single public IP address that the external provider needs for the allow list. The NAT gateway will allow the application to access the service, while keeping the underlying Lambda functions private.
When configuring NAT gateways, you should ensure that the route table associated with the NAT gateway has a route to the internet gateway with a target of the internet gateway. Additionally, you should ensure that the security group associated with the NAT gateway allows outbound traffic from the Lambda functions.
AWS Certified Solutions Architect Professional Official Amazon Text Book [1], page 456
https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Gateway.html

QUESTION 47
A solutions architect has developed a web application that uses an Amazon API Gateway Regional endpoint and an AWS Lambda function. The consumers of the web application are all close to the AWS Region where the application will be deployed. The Lambda function only queries an Amazon Aurora MySQL database. The solutions architect has configured the database to have three read replicas.
During testing, the application does not meet performance requirements. Under high load, the application opens a large number of database connections. The solutions architect must improve the application's performance.
Which actions should the solutions architect take to meet these requirements? (Choose two.)

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

Correct Answer: B, D
Section:
Explanation:
Connect to RDS outside of Lambda handler method to improve performance https://awstut.com/en/2022/04/30/connect-to-rds-outside-of-lambda-handler-method-to-improve-performance-en/
Using RDS Proxy, you can handle unpredictable surges in database traffic. Otherwise, these surges might cause issues due to oversubscribing connections or creating new connections at a fast rate. RDS Proxy establishes a database connection pool and reuses connections in this pool. This approach avoids the memory and CPU overhead of opening a new database connection each time. To protect the database against oversubscription, you can control the number of database connections that are created. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

QUESTION 48
A company is planning to host a web application on AWS and works to load balance the traffic across a group of Amazon EC2 instances. One of the security requirements is to enable end-to-end encryption in transit between the client and the web server.
Which solution will meet this requirement?

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

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

QUESTION 49
A company wants to migrate its data analytics environment from on premises to AWS The environment consists of two simple Node js applications One of the applications collects sensor data and loads it into a MySQL database The other application aggregates the data into reports When the aggregation jobs run. some of the load jobs fail to run correctly
The company must resolve the data loading issue The company also needs the migration to occur without interruptions or changes for the company's customers
What should a solutions architect do to meet these requirements?

A. Set up an Amazon Aurora MySQL database as a replication target for the on-premises database Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica

Set up collection endpomts as AWS Lambda functions behind a Network Load Balancer (NLB). and use Amazon RDS Proxy to wnte to the Aurora MySQL database When the databases are synced disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the NLB.

B. Set up an Amazon Aurora MySQL database Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora Move the aggregation jobs to run against the Aurora MySQL database Set up collection endpomts behind an Application Load Balancer (ALB) as Amazon EC2 instances in an Auto Scaling group When the databases are synced, point the collector DNS record to the ALB Disable the AWS DMS sync task after the cutover from on premises to AWS

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

D. Set up an Amazon Aurora MySQL database Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica Set up collection endpoints as an Amazon Kinesis data stream Use Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database When the databases are synced disable the replication job and restart the Aurora Replica as the primary instance Point the collector DNS record to the Kinesis data stream.

Correct Answer: C
Section:
Explanation:
Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.
Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%

QUESTION 50
A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled.
Which solution will meet these requirements?

A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed key. Use the AWS CLI to re-upload all objects in the S3 bucket. Set an S3 bucket policy to deny unencrypted PutObject requests.

B. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject requests. Use the AWS CLI to re-upload all objects in the S3 bucket.

C. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to automatically encrypt objects on GetObject and PutObject requests.

D. In the S3 bucket properties, change the default encryption to AES-256 with a customer managed key. Attach a policy to deny unencrypted PutObject requests to any entities that access the S3 bucket. Use the AWS CLI to re-upload all objects in the S3 bucket.

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

QUESTION 51
A company is running a web application in the AWS Cloud. The application consists of dynamic content that is created on a set of Amazon EC2 instances. The
EC2 instances run in an Auto Scaling group that is configured as a target group for an Application Load Balancer (ALB).
The company is using an Amazon CloudFront distribution to distribute the application globally. The CloudFront distribution uses the ALB as an origin. The company uses Amazon Route 53 for DNS and has created an A record of www.example.com for the CloudFront distribution.
A solutions architect must configure the application so that itis highly available and fault tolerant.
Which solution meets these requirements?

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

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html
You can set up CloudFront with origin failover for scenarios that require high availability. To get started, you create an origin group with two origins: a primary and a secondary. If the primary origin is unavailable, or returns specific HTTP response status codes that indicate a failure, CloudFront automatically switches to the secondary origin.

QUESTION 52
A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts AWS Site-to-Site VPN connections are configured between ail of the company's global offices and the transit account The company has AWS Config enabled on all of its accounts.
The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices Developers Will reference this list to gain access to applications securely.
Which solution meets these requirements with the LEAST amount of operational overhead?

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

Correct Answer: C
Section:
Explanation:
Customer-managed prefix lists --- Sets of IP address ranges that you define and manage. You can share your prefix list with other AWS accounts, enabling those accounts to reference the prefix list in their own resources.
https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html
a VPC prefix list is created in the transit account with all of the internal IP address ranges, and then shared to all of the other accounts using AWS Resource Access Manager. This allows for central management of the IP address ranges, and eliminates the need for manual updates to security group rules in each account. This solution also allows for compliance checks to be run using AWS Config and for any non-compliant security groups to be automatically remediated.

QUESTION 53
A company runs a new application as a static website in Amazon S3. The company has deployed the application to a production AWS account and uses Amazon CloudFront to deliver the website. The website calls an Amazon API Gateway REST API. An AWS Lambda function backs each API method.
The company wants to create a CSV report every 2 weeks to show each API Lambda function's recommended configured memory, recommended cost, and the price difference between current configurations and the recommendations. The company will store the reports in an S3 bucket.
Which solution will meet these requirements with the LEAST development time?

A. Create a Lambda function that extracts metrics data for each API Lambda function from Amazon CloudWatch Logs for the 2-week penod_ Collate the data into tabular format. Store the data as a _csvfile in an S3 bucket. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.
B. Opt in to AWS Compute Optimizer. Create a Lambda function that calls the ExportLambdaFunctionRecommendatlons operation. Export the _csv file to an S3 bucket. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.

C. Opt in to AWS Compute Optimizer. Set up enhanced infrastructure metrics. Within the Compute Optimizer console, schedule a job to export the Lambda recommendations to a _csvfile_ Store the file in an S3 bucket every 2 weeks.

D. Purchase the AWS Business Support plan for the production account. Opt in to AWS Compute Optimizer for AWS Trusted Advisor checks. In the Trusted Advisor console, schedule a job to export the cost optimization checks to a _csvfile_ Store the file in an S3 bucket every 2 weeks.fs20

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/compute-optimizer/latest/APIReference/API_ExportLambdaFunctionRecommendations.html

QUESTION 54
The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.
Which combination of actions will meet these requirements? (Select THREE.)

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

Correct Answer: A, C, F
Section:
Explanation:
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html
https://aws.amazon.com/premiumsupport/knowledge-center/cost-explorer-analyze-spending-and-usage/
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html https://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html
The best combination of actions to meet the company's requirements is Options A, C, and F.
Option A involves activating the user-defined cost allocation tags that represent the application and the team. This will allow the company to assign costs to different applications or teams, and will allow them to be tracked in the monthly AWS bill.
Option C involves creating a cost category for each application in Billing and Cost Management. This will allow the company to easily identify and compare costs across different applications and teams.
Option F involves enabling Cost Explorer. This will allow the company to view the costs of their AWS resources over the last 12 months and to create forecasts for the next 12 months.
These recommendations are in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that ''You can use cost allocation tags to group your costs by application, team, or other categories'' (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf). Additionally, the book states that ''Cost Explorer enables you to view the costs of your AWS resources over the last 12 months and to create forecasts for the next 12 months'' (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf).

QUESTION 55
An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list. The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets.
How should a solutions architect ensure that the web application can continue to call the third-parly API after the migration?

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

Correct Answer: B

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

QUESTION 56

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

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

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

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

Correct Answer: C
Section:
Explanation:
However A's explanation is incorrect - https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
'SCPs are similar to AWS Identity and Access Management (IAM) permission policies and use almost the same syntax. However, an SCP never grants permissions.'
SCPs alone are not sufficient to granting permissions to the accounts in your organization. No permissions are granted by an SCP. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

QUESTION 57
A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon
Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.
Which solution will meet these requirements?

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

D. Create an image of the instance. Launch a new EC2 instance from the image. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

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

QUESTION 58
A solutions architect needs to copy data from an Amazon S3 bucket m an AWS account to a new S3 bucket in a new AWS account. The solutions architect must implement a solution that uses the AWS CLI.
Which combination of steps will successfully copy the data? (Choose three.)

A. Create a bucket policy to allow the source bucket to list its contents and to put objects and set object ACLs in the destination bucket. Attach the bucket policy to the destination bucket.
B. Create a bucket policy to allow a user In the destination account to list the source bucket's contents and read the source bucket's objects. Attach the bucket policy to the source bucket.
C. Create an IAM policy in the source account. Configure the policy to allow a user In the source account to list contents and get objects In the source bucket, and to list contents, put objects, and set object ACLs in the destination bucket. Attach the policy to the user _
D. Create an IAM policy in the destination account. Configure the policy to allow a user In the destination account to list contents and get objects In the source bucket, and to list contents, put objects, and set objectACLs in the destination bucket. Attach the policy to the user.
E. Run the aws s3 sync command as a user in the source account. Specify' the source and destination buckets to copy the data.
F. Run the aws s3 sync command as a user in the destination account. Specify' the source and destination buckets to copy the data.

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

QUESTION 59
A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.
Which solution will meet these requirements?

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

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-traffic-shifting/
https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html

QUESTION 60
A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files ate uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.examWe.com through the use of Amazon Route 53.

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

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

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html

QUESTION 61
A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.
What should the solutions architect do to meet these requirements?

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

Correct Answer: B
Section:
Explanation:

QUESTION 62
A company is building a solution in the AWS Cloud. Thousands or devices will connect to the solution and send data. Each device needs to be able to send and receive data in real time over the MQTT protocol. Each device must authenticate by using a unique X.509 certificate.
Which solution will meet these requirements with the LEAST operational overhead?

A. Set up AWS loT Core. For each device, create a corresponding Amazon MQ queue and provision a certificate. Connect each device to Amazon MQ.

B. Create a Network Load Balancer (NLB) and configure it with an AWS Lambda authorizer. Run an MQTT broker on Amazon EC2 instances in an Auto Scaling group. Set the Auto Scaling group as the target for the NLB. Connect each device to the NLB.

C. Set up AWS loT Core. For each device, create a corresponding AWS loT thing and provision a certificate. Connect each device to AWS loT Core.

D. Set up an Amazon API Gateway HTTP API and a Network Load Balancer (NLB). Create integration between API Gateway and the NLB. Configure a mutual TLS certificate authorizer on the HTTP API. Run an MQTT broker on an Amazon EC2 instance that the NLB targets. Connect each device to the NLB.

Correct Answer: C
Section:

Explanation:
https://aws.amazon.com/cn/iot-core/faqs/?nc=sn&loc=5&dn=2

QUESTION 63
A company is running several workloads in a single AWS account. A new company policy states that engineers can provision only approved resources and that engineers must use AWS CloudFormation to provision these resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.
What should the solutions architect do to create the solution?

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

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/security-best-practices.html#use-iam-to-control-access https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html

QUESTION 64
A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.
The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.
Which storage strategy is the MOST cost-effective and meets the design requirements?

A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieval. Configure a lifecycle policy to delete data older than 120 days.
B. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scale. Configure the DynamoOB Time to Live (TTL) feature to delete records older than 120 days.
C. Design the application to store each incoming record in a single table in an Amazon RDS MySQL database. Run a nightly cron job that executes a query to delete any records older than 120 days.
D. Design the application to batch incoming records before writing them to an Amazon S3 bucket. Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the data. Configure a lifecycle policy to delete the data after 120 days.

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

QUESTION 65
A retail company is hosting an ecommerce website on AWS across multiple AWS Regions. The company wants the website to be operational at all times for online purchases. The website stores data in an Amazon RDS for MySQL DB instance.
Which solution will provide the HIGHEST availability for the database?

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

D. Configure read replicas on Amazon RDS. In the case of disruption, promote a cross-Region and read replica to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.

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

QUESTION 66
A company has its cloud infrastructure on AWS A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts
What should the solutions architect do to meet these requirements?

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

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-organization/
AWS Organizations allows the management of multiple AWS accounts as a single entity and AWS CloudFormation StackSets allows creating, updating, and deleting stacks across multiple accounts and regions in an organization. This solution allows creating a single CloudFormation template that can be deployed across multiple accounts and regions, and also allows for the management of access and permissions for the different accounts through the use of IAM roles and policies in the management account.

QUESTION 67
A company plans to refactor a monolithic application into a modern application designed deployed or AWS. The CLCD pipeline needs to be upgraded to support the modem design for the application with the following requirements
* It should allow changes to be released several times every hour.
* It should be able to roll back the changes as quickly as possible
Which design will meet these requirements?

A. Deploy a Cl-CD pipeline that incorporates AMIs to contain the application and their configurations Deploy the application by replacing Amazon EC2 instances
B. Specify AWS Elastic Beanstak to sage in a secondary environment as the deployment target for the CI/CD pipeline of the application. To deploy swap the staging and production environment URLs.
C. Use AWS Systems Manager to re-provision the infrastructure for each deployment Update the Amazon EC2 user data to pull the latest code art-fact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment
D. Roll out At application updates as pan of an Auto Scaling event using prebuilt AMIs. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Correct Answer: B
Section:
Explanation:
It is the fastest when it comes to rollback and deploying changes every hour

QUESTION 68
A company has an application that runs on Amazon EC2 instances. A solutions architect is designing VPC infrastructure in an AWS Region where the application needs to access an Amazon Aurora DB cluster. The EC2 instances are all associated with the same security group. The DB cluster is associated with its own security group.

The solutions architect needs to add rules to the security groups to provide the application with least privilege access to the DB cluster.
Which combination of steps will meet these requirements? (Select TWO.)

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

Correct Answer: B, C
Section:
Explanation:
Outbound rule to the EC2 SG with DB SG as destination
Inbound rule to the DB SG with EC2 SG as source

QUESTION 69
A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations lo manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.
Which solution is the MOST cost-effective way to meet these requirements?

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

Correct Answer: B
Section:
Explanation:
Configure AWS Budgets in the organizations master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organizations master account to create monthly reports for each business unit.
https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Budgets%20gives%20you%20the,below%20the%20threshold%20you%20define

QUESTION 70
A video processing company wants to build a machine learning (ML) model by using 600 TB of compressed data that is stored as thousands of files in the company's on-premises network attached storage system. The company does not have the necessary compute resources on premises for ML experiments and wants to use AWS.
The company needs to complete the data transfer to AWS within 3 weeks. The data transfer will be a one-time transfer. The data must be encrypted in transit. The measured upload speed of the company's internet connection is 100 Mbps, and multiple departments share the connection.
Which solution will meet these requirements MOST cost-effectively?

A. Order several AWS Snowball Edge Storage Optimized devices by using the AWS Management Console. Configure the devices with a destination S3 bucket. Copy the data to the devices. Ship the devices back to AWS.
B. Set up a 10 Gbps AWS Direct Connect connection between the company location and the nearest AWS Region. Transfer the data over a VPN connection into the Region to store the data in Amazon S3.
C. Create a VPN connection between the on-premises network storage and the nearest AWS Region. Transfer the data over the VPN connection.
D. Deploy an AWS Storage Gateway file gateway on premises. Configure the file gateway with a destination S3 bucket. Copy the data to the file gateway.

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

QUESTION 71
A company has migrated its forms-processing application to AWS. When users interact with the application, they upload scanned forms as files through a web application. A database stores user metadata and references to files that are stored in Amazon S3. The web application runs on Amazon EC2 instances and an Amazon RDS for PostgreSQL database.
When forms are uploaded, the application sends notifications to a team through Amazon Simple Notification Service (Amazon SNS). A team member then logs in and processes each form. The team member performs data validation on the form and extracts relevant data before entering the information into another system that uses an API.
A solutions architect needs to automate the manual processing of the forms. The solution must provide accurate form extraction, minimize time to market, and minimize long-term operational overhead.
Which solution will meet these requirements?

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

Correct Answer: D
Section:
Explanation:
Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API. This solution meets the requirements of accurate form extraction, minimal time to market, and minimal long-term operational overhead. Amazon Textract and Amazon Comprehend are fully managed and serverless services that can perform OCR and extract relevant data from the forms, which eliminates the need to develop custom libraries or train and host models. Using AWS Step Functions and Lambda allows for easy automation of the process and the ability to scale as needed.

QUESTION 72
A company is refactoring its on-premises order-processing platform in the AWS Cloud. The platform includes a web front end that is hosted on a fleet of VMs RabbitMQ to connect the front end to the backend, and a Kubernetes cluster to run a containerized backend system to process the orders. The company does not want to make any major changes to the application
Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: A

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

QUESTION 73
A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.
The solutions architect created the following IAM policy and attached it to an IAM role:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DownloadUpload",
            "Action": [
                "s3:GetObject",
                "s3:GetObjectVersion",
                "s3:PutObject",
                "s3:PutObjectAcl"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:s3:::BucketName/*"
        },
        {
            "Sid": "KMSAccess",
            "Action": [
                "kms:Decrypt",
                "kms:Encrypt"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:kms:Region:Account:key/Key ID"
        }
    ]
}
```

During tests, me solutions architect was able to successfully get existing test objects m the S3 bucket However, attempts to upload a new object resulted in an error message. The error message stated that me action was forbidden.
Which action must me solutions architect add to the IAM policy to meet all the requirements?

A. Kms:GenerateDataKey
B. KmsGetKeyPolpcy
C. kmsGetPubKKey
D. kms:SKjn

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-access-denied-error-kms/
'An error occurred (AccessDenied) when calling the PutObject operation: Access Denied' This error message indicates that your IAM user or role needs permission for the kms:GenerateDataKey action.

QUESTION 74
A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.
How should a solutions architect configure the web ACLs to meet these requirements?

A. Set the action of the web ACL rules to Count. Enable AWS WAF logging Analyze the requests for false positives Modify the rules to avoid any false positive Over time change the action of the web ACL rules from Count to Block.
B. Use only rate-based rules in the web ACLs. and set the throttle limit as high as possible Temporarily block all requests that exceed the limit. Define nested rules to narrow the scope of the rate tracking.

C. Set the action o' the web ACL rules to Block. Use only AWS managed rule groups in the web ACLs Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.
D. Use only custom rule groups in the web ACLs. and set the action to Allow Enable AWS WAF logging Analyze the requests tor false positives Modify the rules to avoid any false positive Over time, change the action of the web ACL rules from Allow to Block.

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/waf-analyze-count-action-rules/

QUESTION 75
A company has an organization that has many AWS accounts in AWS Organizations. A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.
The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to and from the company's on-premises network.
Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.
The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS account. Deploy an AWS Lambda function in each AWS account. Configure the Lambda function to run every time an SNS topic receives a message. Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account. Instruct the security team to distribute changes by publishing messages to its SNS topic.
B. Create new customer-managed prefix lists in each AWS account within the organization. Populate the prefix lists in each account with all internal CIDR ranges. Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups. Instruct the security team to share updates with each AWS account owner.
C. Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.
D. Create an IAM role in each account in the organization. Grant permissions to update security groups. Deploy an AWS Lambda function in the security team's AWS account. Configure the Lambda function to take a list of internal IP addresses as input, assume a role in each organization account, and add the list of IP addresses to the security groups in each account.

Correct Answer: C
Section:
Explanation:
Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups. This solution meets the requirements with the least amount of operational overhead as it requires the security team to create and maintain a single customer-managed prefix list, and share it with the organization using AWS Resource Access Manager. The owners of each AWS account are then responsible for allowing the prefix list in their security groups, which eliminates the need for the security team to manually notify each account owner when changes are made. This solution also eliminates the need for a separate AWS Lambda function in each account, reducing the overall complexity of the solution.

QUESTION 76
A company has introduced a new policy that allows employees to work remotely from their homes if they connect by using a VPN The company Is hosting Internal applications with VPCs in multiple AWS accounts Currently the applications are accessible from the company's on-premises office network through an AWS Site-to-Site VPN connection The VPC in the company's main AWS account has peering connections established with VPCs in other AWS accounts.
A solutions architect must design a scalable AWS Client VPN solution for employees to use while they work from home
What is the MOST cost-effective solution that meets these requirements?

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

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/scenario-peered.html

QUESTION 77
A company is running an application in the AWS Cloud. Recent application metrics show inconsistent response times and a significant increase in error rates. Calls to third-party services are causing the delays. Currently, the application calls third-party services synchronously by directly invoking an AWS Lambda function.
A solutions architect needs to decouple the third-party service calls and ensure that all the calls are eventually completed.
Which solution will meet these requirements?

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

Correct Answer: A
Section:
Explanation:
Using an SQS queue to store events and invoke the Lambda function will decouple the third-party service calls and ensure that all the calls are eventually completed. SQS allows you to store messages in a queue and process them asynchronously, which eliminates the need for the application to wait for a response from the third-party service. The messages will be stored in the SQS queue until they are processed by the Lambda function, even if the Lambda function is currently unavailable or busy. This will ensure that all the calls are eventually completed, even if there are delays or errors.
AWS Step Functions state machines can also be used to pass events to the Lambda function, but it would require additional management and configuration to set up the state machine, which would increase operational overhead.
Amazon EventBridge rule can also be used to pass events to the Lambda function, but it would not provide the same level of decoupling and reliability as SQS.
Using Amazon Simple Notification Service (Amazon SNS) topic to store events and Invoke the Lambda function, is similar to SQS, but SNS is a publish-subscribe messaging service and SQS is a queue service. SNS is used for sending messages to multiple recipients, SQS is used for sending messages to a single recipient, so SQS is more appropriate for this use case.
AWS SQS
AWS Step Functions
AWS EventBridge
AWS SNS

QUESTION 78
A company is running applications on AWS in a multi-account environment. The company's sales team and marketing team use separate AWS accounts in AWS Organizations.
The sales team stores petabytes of data in an Amazon S3 bucket. The marketing team uses Amazon QuickSight for data visualizations. The marketing team needs access to data that the sates team stores in the S3 bucket. The company has encrypted the S3 bucket with an AWS Key Management Service (AWS KMS) key. The marketing team has already created the IAM service role for QuickSight to provide QuickSight access in the marketing AWS account. The company needs a solution that will provide secure access to the data in the S3 bucket across AWS accounts.
Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: D

Section:
Explanation:
Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight role, to create a trust relationship with the new IAM role in the sales account.
This approach is the most secure way to grant cross-account access to the data in the S3 bucket while minimizing operational overhead. By creating an IAM role in the sales account, the marketing team can assume the role in their own account, and have access to the S3 bucket. And updating the QuickSight role, to create a trust relationship with the new IAM role in the sales account will grant the marketing team to access the data in the S3 bucket and use it for data visualization using QuickSight.
AWS Resource Access Manager (AWS RAM) also allows sharing of resources between accounts, but it would require additional management and configuration to set up the sharing, which would increase operational overhead.
Using S3 replication would also replicate the data to the marketing account, but it would not provide the marketing team access to the original data, and also it would increase operational overhead with managing the replication process.
IAM roles and policies, KMS grants and trust relationships are a powerful combination for managing cross-account access in a secure and efficient manner.
AWS IAM Roles
AWS KMS - Key Grants
AWS RAM

QUESTION 79
A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a
Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS.
Which solution will meet these requirements?

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

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/dms/schema-conversion-tool/
AWS Schema Conversion Tool (SCT) can automatically convert the database schema from Microsoft SQL Server to Amazon RDS for MySQL. This allows for a smooth transition of the database schema without any manual intervention. AWS DMS can then be used to migrate the data from the on-premises databases to the newly created Amazon RDS for MySQL instance. This service can perform a one-time migration of the data or can set up ongoing replication of data changes to keep the on-premises and AWS databases in sync.

QUESTION 80
A publishing company's design team updates the icons and other static assets that an ecommerce web application uses. The company serves the icons and assets from an Amazon S3 bucket that is hosted in the company's production account. The company also uses a development account that members of the design team can access.
After the design team tests the static assets in the development account, the design team needs to load the assets into the S3 bucket in the production account. A solutions architect must provide the design team with access to the production account without exposing other parts of the web application to the risk of unwanted changes.
Which combination of steps will meet these requirements? (Select THREE.)

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

Correct Answer: A, C, E
Section:
Explanation:
A) In the production account, create a new IAM policy that allows read and write access to the S3 bucket. The policy grants the necessary permissions to access the assets in the production S3 bucket.
C) In the production account, create a role. Attach the new policy to the role. Define the development account as a trusted entity. By creating a role and attaching the policy, and then defining the development account as a trusted entity, the development account can assume the role and access the production S3 bucket with the read and write permissions.
E) In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account. The IAM policy attached to the group allows the design team members to assume the role created in the production account, thereby giving them access to the production S3 bucket.
Step 1: Create a role in the Production Account; create the role in the Production account and specify the Development account as a trusted entity. You also limit the role permissions to only read and write access to the productionapp bucket. Anyone granted permission to use the role can read and write to the productionapp bucket. Step 2: Grant access to the role Sign in as an administrator in the Development account and allow the AssumeRole action on the UpdateApp role in the Production account. So, recap, production account you create the policy for S3, and you set development account as a trusted entity. Then on the development account you allow the sts:assumeRole action on the role in production account. https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

QUESTION 81
A company developed a pilot application by using AWS Elastic Beanstalk and Java. To save costs during development, the company's development team deployed the application into a single-instance environment. Recent tests indicate that the application consumes more CPU than expected. CPU utilization is regularly greater than 85%, which causes some performance bottlenecks.
A solutions architect must mitigate the performance issues before the company launches the application to production.
Which solution will meet these requirements with the LEAST operational overhead?

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

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

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

Correct Answer: C
Section:
Explanation:
This solution will meet the requirements with the least operational overhead because it allows the company to modify the existing environment's capacity configuration, so it becomes a load-balanced environment type. By selecting all availability zones, the company can ensure that the application is running in multiple availability zones, which can help to improve the availability and scalability of the application. The company can also add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes, which can help to mitigate the performance issues. This solution does not require creating new Elastic Beanstalk environments or rebuilding the existing one, which reduces the operational overhead.
You can refer to the AWS Elastic Beanstalk documentation for more information on how to use this service: https://aws.amazon.com/elasticbeanstalk/ You can refer to the AWS documentation for more information on how to use autoscaling: https://aws.amazon.com/autoscaling/

QUESTION 82
A finance company is running its business-critical application on current-generation Linux EC2 instances The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.
Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

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

Correct Answer: B

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

QUESTION 83
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 10 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 documentation1, AWS App2Container (A2C) is a command line tool for migratingand modernizing Java and .NET web applications into container format. AWS A2C analyzes and builds aninventory 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 anddeploy them on Amazon ECS or Amazon EKS.Option A meets the requirements of the scenario because it allows you to migrate your existing Javaapplication to AWS and minimize the administrative overhead to maintain the servers. You can use AWSA2C 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 registryservice. You can use AWS Fargate as the launch type for your Amazon ECS cluster, which is a serverlesscompute engine that eliminates the need to provision and manage servers for your containers. You cangrant the ECS task execution role permission to access the ECR image repository, which allows your tasksto pull images from ECR. You can configure Amazon ECS to use an ALB, which is a load balancer thatdistributes traffic across multiple targets in multiple Availability Zones using HTTP or HTTPS protocols.You can use the ALB to interact with your application.

QUESTION 84
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: D
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 85
A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS)

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

A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbooks. Copy the snapshots to the secondary Region. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.

B. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volumes. In the event of a failure, launch 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.

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 volumes1. Configuring the backup task to copy the backups to a vault in the secondary Region will enable maintaining backups in a separate Region1. In the event of a failure, launching the CloudFormation template will enable deploying the network configuration in the secondary Region2. Restoring the instance volumes and configurations from the backup vault will enable recovering the EC2 instances and their data1. Transferring usage to the secondary Region will enable resuming operations2.

QUESTION 86
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 fi-nance 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 con-trol). Establish trust with the marketing team's account. In the mar-keting team's account, create an IAM role that has permissions to as-sume the IAM role in the finance team's account.

C. Create a resource-based IAM policy that includes conditions for spe-cific 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 Dyna-moDB 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 names1. By creating a resource-based policy that allows access to only

specific attributes of the DynamoDB table and attaching it to the table, the company can restrict access to confidential data. By creating an IAM role in the marketing team's account that has permissions to access the DynamoDB table in the finance team's account, the company can enable cross-account access.
The other options are not correct because:
Creating an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table would not work because SCPs are policies that you can use with AWS Organizations to manage permissions in your organization's accounts.SCPs do not grant permissions; instead, they specify the maximum permissions that identities in an account can have2. SCPs cannot be used to specify fine-grained access control for DynamoDB items and attributes.
Creating an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes and establishing trust with the marketing team's account would not work because IAM roles are identities that you can create in your account that have specific permissions.You can use an IAM role to delegate access to users, applications, or services that don't normally have access to your AWS resources3. However, creating an IAM role in the finance team's account would not restrict access to specific attributes of the DynamoDB table; it would only allow cross-account access. The company would still need a resource-based policy attached to the table to enforce fine-grained access control.
Creating an IAM role in the finance team's account to access the DynamoDB table and using an IAM permissions boundary to limit the access to the specific attributes would not work because IAM permissions boundaries are policies that you use to delegate permissions management to other users.You can use permissions boundaries to limit the maximum permissions that an identity-based policy can grant to an IAM entity (user or role)4. Permissions boundaries cannot be used to specify fine-grained access control for DynamoDB items and attributes.
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/specifying-conditions.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

QUESTION 87
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-I 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-l 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 users1. AWS Global Accelerator is a service that uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies1. Configuring the ALB as the origin will enable connecting the accelerator to the ALB that exposes the APIs2. AWS Global Accelerator supports non-standard REST methods such as LINK, UNLINK, LOCK, and UNLOCK3.

QUESTION 88
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 o 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 89
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 90
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)1. Creating a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses will enable enforcing encryption for all requests to the bucket2. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS3.

QUESTION 91
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 accounts1. Using AWS Firewall Manager to deploy AWS WAF rules in all accounts for all CloudFront distributions will enable providing baseline protection for the OWASP top 10 web application vulnerabilities2. AWS Firewall Manager supports AWS WAF rules that can help protect against common web exploits such as SQL injection and cross-site scripting3. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS.

QUESTION 92
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 servers1.
* Use AWS AppSync to create a scalable GraphQL API that connects to your database and other data sources1.
* Use Amazon SQS to decouple and scale your order processing microservices1.
* Use AWS Lambda to run code for your business logic without provisioning or managing servers1.
* Use an Amazon SQS dead-letter queue to retain messages that can't be processed by your Lambda function1.

QUESTION 93
A company needs to build a disaster recovery (DR) solution for its ecommerce website. The web application is hosted on a fleet of t3.Iarge 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 R TO 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 minutes1. AWS DRS continuously replicates data from the source servers to a staging area subnet in the DR Region, where it uses low-cost storage and minimal compute resources to maintain ongoing replication.In the event of a disaster, AWS DRS automatically converts the servers to boot and run natively on AWS and launches recovery instances on AWS within minutes2. By using AWS DRS, the company can save costs by removing idle recovery site resources and paying for the full disaster recovery site only when needed.By creating a cross-Region read replica for the DB instance, the company can have a standby copy of its primary database in a different AWS Region3.By using infrastructure as code (IaC), the company can provision the new infrastructure in the DR Region in an automated and consistent way4. By using an Amazon Route 53 failover routing policy, the company can route traffic to a resource that is healthy or to another resource when the first resource becomes unavailable.
The other options are not correct because:
Using AWS Backup to create cross-Region backups for the EC2 instances and the DB instance would not meet the RPO and RTO requirements. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. You can use AWS Backup to back up your application data across AWS services in your account and across accounts. However, AWS Backup does not provide continuous replication or fast recovery; it creates backups at scheduled intervals and requires manual restoration. Creating backups every 30 seconds would also incur high costs and network bandwidth.
Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data.
Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying dat
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://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 94
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 Region1. By creating a cross-Region read replica for the RDS database in the new Region, the company can have a standby copy of its primary database that can serve read-only traffic from users in Europe. A latency-based routing policy is a feature that enables you to route traffic based on the latency between your users and your resources.You can use latency-based routing to route traffic to the resource that provides the best latency2. By changing the Route 53 record to latency-based routing, the company can minimize latency for users who download reports by connecting them to the API Gateway API in the Region that provides the best response time.
The other options are not correct because:
Using AWS Database Migration Service (AWS DMS) to replicate the primary database in the original Region to the database in the new Region would not be as cost-effective or simple as using a cross-Region read replica. AWS DMS is a service that enables you to migrate relational databases, data warehouses, NoSQL databases, and other types of data stores.You can use AWS DMS to perform one-time migrations or continuous data replication with high availability and consolidate databases into a petabyte-scale data warehouse3. However, AWS DMS requires more configuration and management than creating a cross-Region read replica, which is fully managed by Amazon RDS. AWS DMS also incurs additional charges for replication instances and tasks.
Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery or minimizing latency. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data from an RDS database.
Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery or minimizing latency. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data from an RDS database.
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 95

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 I-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 discount1. Savings Plans are flexible pricing models that offer low prices on EC2 instances, Fargate, and Lambda usage, in exchange for a commitment to a consistent amount of usage (measured in $/hour) for a one- or three-year term. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any EC2 instance regardless of Region, instance family, operating system, or tenancy, including those that are part of EMR, ECS, or EKS clusters, or launched by Fargate or Lambd
a.EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings2. By purchasing the same Reserved Instances for an additional 3-year term with All Upfront payment, the company can lock in the lowest possible price for its EC2 instances that run continuously for 3 years. By purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account, the company can benefit from additional discounts on any other compute usage across its member accounts.
The other options are not correct because:
Purchasing a 1-year Compute Savings Plan with No Upfront payment in each member account would not provide as much cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. A 1-year term offers lower discounts than a 3-year term, and a No Upfront payment option offers lower discounts than an All Upfront payment option. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.
Purchasing a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region would not provide as much cost savings as purchasing Reserved Instances for an additional 3-year term with All Upfront payment. An EC2 Instance Savings Plan offers lower discounts than Reserved Instances for the same instance family and Region. Also, a No Upfront payment option offers lower discounts than an All Upfront payment option.
Purchasing a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account would not provide as much flexibility or cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. An EC2 Instance Savings Plan applies only to a specific instance family within a Region and does not cover Fargate or Lambda usage. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.
https://aws.amazon.com/ec2/pricing/reserved-instances/
https://aws.amazon.com/savingsplans/

QUESTION 96
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 requests1. By creating a second target group and deploying the new logic to EC2 instances in this new target group, the company can have two versions of its business logic running in parallel. By updating the ALB listener rule to use weighted target groups, the company can control how much traffic is sent to each version. By configuring ALB target group stickiness, the company can ensure that a customer uses the same version of the business logic during the testing window.Target group stickiness is a feature that enables you to bind a user's session to a specific target within a target group for the duration of the session2.
The other options are not correct because:
Creating a second ALB and deploying the new logic to a set of EC2 instances in a new Auto Scaling group would not be as cost-effective or simple as using weighted target groups. A second ALB would incur additional charges and require more configuration and management. Updating the Route 53 record to use weighted routing would not ensure that a customer uses the same version of the business logic during the testing window, as DNS caching could affect how requests are routed.
Creating a new launch configuration for the Auto Scaling group and replacing it on the Auto Scaling group would not allow for gradual traffic shifting between versions. A launch configuration is a template that an Auto Scaling group uses to launch EC2 instances.You can specify information such as the AMI ID, instance type, key pair, security groups, and block device mapping for your instances3. However, replacing the launch configuration on an Auto Scaling group would affect all instances in that group, not just 10% of customers.
Creating a second Auto Scaling group and changing the ALB routing algorithm to least outstanding requests (LOR) would not allow for controlled traffic shifting between versions. A second Auto Scaling group would require more configuration and management. The LOR routing algorithm is a feature that enables you to route traffic based on how quickly targets respond to requests.The load balancer selects a target from the target group with the fewest outstanding requests4. However, this algorithm does not take into account customer sessions or weights.
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 97
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 needed1. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs. AWS Amplify offers the following features and benefits:
Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.
Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.
Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.
Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.
Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.
By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a 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 data2.By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users3.
The other options are not correct because:
Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.
Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.
Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.
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 98
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 06 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 needed1. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs. AWS Amplify offers the following features and benefits:
Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.
Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.
Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.
Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.
Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.
By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a 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 data2.By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users3.
The other options are not correct because:
Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.
Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.
Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.
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 99
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 Replic a. 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 100

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 files1. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances.A Multi-AZ file system replicates data across multiple Availability Zones (AZs) within a Region, providing high availability and durability2. By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput.Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system's size or burst credit balance3. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of less than 1 hour.
The other options are not correct because:
Deploying a new Amazon FSx for Lustre file system would not provide a single location for updates to application data for all instances. Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance storage for compute workloads. However, it does not support concurrent write access from multiple instances. Using AWS Backup to back up the file system to the DR Region would not provide real-time replication of dat
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 101
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 blocked1. By creating a new AWS WAF web ACL, the company can create a collection of rules that define the conditions for allowing or blocking web requests. By adding a new rule that blocks requests that match the SQL database rule group, the company can prevent SQL injection attacks from reaching the ECS API service.The SQL database rule group is a managed rule group provided by AWS that contains rules to protect against common SQL injection attack patterns2. By setting the web ACL to allow all other traffic that does not match those rules, the company can ensure that legitimate traffic can access the API service. By attaching the web ACL to the ALB in front of the ECS tasks, the company can apply the web security rules to all requests that are forwarded by the load balancer.
The other options are not correct because:
Creating a new AWS WAF Bot Control implementation would not prevent SQL injection attacks from reaching the ECS API service. AWS WAF Bot Control is a feature that gives you visibility and control over common and pervasive bot traffic that can consume excess resources, skew metrics, cause downtime, or perform other undesired activities.However, it does not protect against SQL injection attacks, which are malicious attempts to execute unauthorized SQL statements against your database3.
Creating a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks would not prevent SQL injection attacks from reaching the ECS API service. Monitoring mode is a feature that enables you to evaluate how your rules would perform without actually blocking any requests.However, this mode does not provide any protection against attacks, as it only logs and counts requests that match your rules4.
Creating a new AWS WAF web ACL and creating a new empty IP set in AWS WAF would not prevent SQL injection attacks from reaching the ECS API service. An IP set is a feature that enables you to specify a list of IP addresses or CIDR blocks that you want to allow or block based on their source IP address. However, this approach would not be effective or efficient against SQL injection attacks, as it would require constantly updating the IP set with new IP addresses of attackers, and it would not block attackers who use proxies or VPNs.
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 102
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 environment1. By creating an Amazon CloudFront distribution with the ALB as the origin, the company can improve the performance and availability of its application by caching static content at edge locations closer to end users. By adding a custom header and random value on the CloudFront domain, the company can prevent direct access to the ALB and ensure that only requests from CloudFront are forwarded to the ECS tasks. By configuring the ALB to conditionally forward traffic if the header and value match, the company can implement origin access identity (OAI) for its ALB origin.OAI is a feature that enables you to restrict access to your content by requiring users to access your content through CloudFront URLs2. By deploying an AWS WAF web ACL that includes an appropriate rule group, the company can prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack. AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications.You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked3. By associating the web ACL with the Amazon CloudFront distribution, the company can apply the web security rules to all requests that are forwarded by CloudFront.
The other options are not correct because:
Deploying the application in two AWS Regions and configuring Amazon Route 53 to route to both Regions with equal weight would not prevent attacks or ensure business continuity.Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service that routes end users to Internet applications by translating names like www.example.com into numeric IP addresses4. However, routing traffic to multiple Regions would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and AWS WAF.
Configuring auto scaling for Amazon ECS tasks and creating a DynamoDB Accelerator (DAX) cluster would not prevent attacks or ensure business continuity. Auto scaling is a feature that enables you to automatically adjust your ECS tasks based on demand or a schedule. DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement. However, these features would not protect against attacks or provide failover in case of an outage. They would also increase operational complexity and costs compared to using CloudFront and AWS WAF.
Configuring Amazon ElastiCache to reduce overhead on DynamoDB would not prevent attacks or ensure business continuity. Amazon ElastiCache is a fully managed in-memory data store service that makes it easy to deploy, operate, and scale popular open-source compatible in-memory data stores. However, this service would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and AWS WAF.
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 103
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 I 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-month1. Therefore, moving the backups to cold storage as soon as possible will reduce the storage

cost.However, cold storage backups must be retained for a minimum of 90 days2, otherwise they incur a pro-rated charge equal to the storage charge for the remaining days1. Therefore, setting the backup retention period to 30 days will incur a penalty of 60 days of cold storage cost for each backup deleted. This penalty will still be lower than keeping the backups in warm storage for 7 days and then in cold storage for 83 days, which is the current configuration. Therefore, option A is the most cost-effective solution.

QUESTION 104
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 sources1.AWS AppSync supports GraphQL subscriptions to perform real-time operations and can push data to clients that choose to listen to specific events from the backend1.AWS AppSync uses WebSockets to establish and maintain a secure connection between the clients and the API endpoint2. Therefore, using AWS AppSync and leveraging WebSockets is a suitable design to reduce comment latency and improve user experience.

QUESTION 105
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 practices1.One of the Trusted Advisor checks is ''Low Utilization Amazon EC2 Instances'', which identifies EC2 instances that appear to be underutilized based on CPU, network I/O, and disk I/O metrics1. This check can help users optimize the cost and size of their EC2 instances by recommending smaller or more appropriate instance types.
AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of AWS resources and generates optimization recommendations to reduce the cost and improve the performance of workloads2.Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWS Fargate, and Lambda functions2.For EC2 instances, Compute Optimizer evaluates the vCPUs, memory, storage, and other specifications, as well as the CPU utilization, network in and out, disk read and write, and other utilization metrics of currently running instances3. It then recommends optimal instance types based on price-performance 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 plan1.
Option C is incorrect because installing the Amazon CloudWatch agent and configuring memory metric collection on the EC2 instances will not provide any optimization recommendations by itself.However, memory metrics can be used by Compute Optimizer to enhance its recommendations if enabled3.
Option E is incorrect because creating an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest will not help with cost-optimization and sizing of EC2 instances.Savings Plans are a flexible pricing model that offer lower prices on Amazon EC2 usage in exchange for a commitment to a consistent amount of usage for a 1- or 3-year term4. Savings Plans do not affect the configuration or utilization of EC2 instances.

QUESTION 106
A company that provides image storage services wants to deploy a customer-lacing 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 deletion1. Users can configure S3 Event Notifications to invoke an AWS Lambda function when a user stores an image in the bucket.Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources2. The Lambda function can resize the image in place and store the original file in the same S3 bucket. This way, the solution can handle significant variance in demand and be reliable at enterprise scale.The solution can also rerun processing jobs in the event of failure by using the retry and dead-letter queue features of Lambda2.
S3 Lifecycle is a feature that allows users to manage their objects so that they are stored cost-effectively throughout their lifecycle3. Users can create an S3 Lifecycle policy to move all stored images to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months.S3 Standard-IA is a storage class designed for data that is accessed less frequently, but requires rapid access when needed4. It offers a lower storage cost than S3 Standard, but charges a retrieval fee. Therefore, moving the images to S3 Standard-IA after 6 months can reduce the storage cost for the solution.
Option A is incorrect because using AWS Step Functions to process the S3 event that occurs when a user stores an image is not necessary or cost-effective. AWS Step Functions is a service that lets users coordinate multiple AWS services into serverless workflows. However, for this use case, a single Lambda function can handle the image resizing task without needing Step Functions.
Option B is incorrect because using Amazon EventBridge to process the S3 event that occurs when a user uploads an image is not necessary or cost-effective. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. However, for this use case, S3 Event Notifications can directly invoke the Lambda function without needing EventBridge.
Option D is incorrect because using Amazon Simple Queue Service (Amazon SQS) to process the S3 event that occurs when a user stores an image is not necessary or cost-effective. Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices, distributed systems, and serverless applications. However, for this use case, S3 Event Notifications can directly invoke the Lambda function without needing SQS. Moreover, storing the resized file in an S3 bucket that uses S3 Standard-IA will incur a retrieval fee every time the file is accessed, which may not be cost-effective for frequently accessed files.

QUESTION 107
A company is running a web application in a VPC. The web application runs on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is using AWS WAF.
An external customer needs to connect to the web application. The company must provide IP addresses to all external customers.
Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html Option A is wrong. AWS WAF does not support associating with NLB. https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html Option B is wrong. An ALB does not support an Elastic IP address. https://aws.amazon.com/elasticloadbalancing/features/

QUESTION 108
A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.
Which combination of steps will meet these requirements? (Choose three.)

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

Correct Answer: C, D, F
Section:
Explanation:
https://docs.aws.amazon.com/controltower/latest/userguide/controls.html https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#ebs-enable-encryption AWS is now transitioning the previous term 'guardrail' new term 'control'.

QUESTION 109
A company is running a critical stateful web application on two Linux Amazon EC2 instances behind an Application Load Balancer (ALB) with an Amazon RDS for MySQL database The company hosts the DNS records for the application in Amazon Route 53 A solutions architect must recommend a solution to improve the resiliency of the application
The solution must meet the following objectives:
* Application tier RPO of 2 minutes. RTO of 30 minutes
* Database tier RPO of 5 minutes RTO of 30 minutes
The company does not want to make significant changes to the existing application architecture The company must ensure optimal latency after a failover
Which solution will meet these requirements?

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

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

QUESTION 110
A company uses an AWS CodeCommit repository The company must store a backup copy of the data that is in the repository in a second AWS Region
Which solution will meet these requirements?

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

Correct Answer: B
Section:
Explanation:
AWS Backup is a fully managed service that makes it easy to centralize and automate the creation, retention, and restoration of backups across AWS services. It provides a way to schedule automatic backups for CodeCommit repositories on an hourly basis. Additionally, it also supports cross-Region replication, which allows you to copy the backups to a second Region for disaster recovery.
By using AWS Backup, the company can set up an automatic and regular backup schedule for the CodeCommit repository, ensuring that the data is regularly backed up and stored in a second Region. This can provide a way to recover quickly from any disaster event that might occur.
AWS Backup documentation: https://aws.amazon.com/backup/
AWS Backup for AWS CodeCommit documentation: https://aws.amazon.com/about-aws/whats-new/2020/07/aws-backup-now-supports-aws-codecommit-repositories/

QUESTION 111
A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company's marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.
Which solution will meet these requirements with the LEAST operational overhead?

A. Instruct each business unit to add a unique secondary CIDR range to the business unit's VPC. Peer the VPCs and use a private NAT gateway in the secondary range to route traffic to the marketing team.
B. Create an Amazon EC2 instance to serve as a virtual appliance in the marketing account's VPC. Create an AWS Site-to-Site VPN connection between the marketing team and each business unit's VPC. Perform NAT where necessary.
C. Create an AWS PrivateLink endpoint service to share the marketing application. Grant permission to specific AWS accounts to connect to the service. Create interface VPC endpoints in other accounts to access the application by using private IP addresses.
D. Create a Network Load Balancer (NLB) in front of the marketing application in a private subnet. Create an API Gateway API. Use the Amazon API Gateway private integration to connect the API to the NLB. Activate IAM authorization for the API. Grant access to the accounts of the other business units.

Correct Answer: C
Section:
Explanation:
With AWS PrivateLink, the marketing team can create an endpoint service to share their internal application with other accounts securely using private IP addresses. They can grant permission to specific AWS accounts to connect to the service and create interface VPC endpoints in the other accounts to access the application by using private IP addresses. This option does not require any changes to the network of the other business units, and it does not require peering or NATing. This solution is both scalable and secure.
https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/

QUESTION 112
A company needs to audit the security posture of a newly acquired AWS account. The company's data security team requires a notification only when an Amazon S3 bucket becomes publicly exposed. The company has already established an Amazon Simple Notification Service (Amazon SNS) topic that has the data security team's email address subscribed.
Which solution will meet these requirements?

A. Create an S3 event notification on all S3 buckets for the isPublic event. Select the SNS topic as the target for the event notifications.
B. Create an analyzer in AWS Identity and Access Management Access Analyzer. Create an Amazon EventBridge rule for the event type ''Access Analyzer Finding'' with a filter for ''isPublic: true.'' Select the SNS topic as the

EventBridge rule target.

C. Create an Amazon EventBridge rule for the event type ''Bucket-Level API Call via CloudTrail'' with a filter for ''PutBucketPolicy.'' Select the SNS topic as the EventBridge rule target.

D. Activate AWS Config and add the cloudtrail-s3-dataevents-enabled rule. Create an Amazon EventBridge rule for the event type ''Config Rules Re-evaluation Status'' with a filter for ''NON_COMPLIANT.'' Select the SNS topic as the EventBridge rule target.

Correct Answer: B
Section:
Explanation:
Access Analyzer is to assess the access policy. https://docs.aws.amazon.com/ja_jp/AmazonS3/latest/userguide/access-control-block-public-access.html

QUESTION 113
A solutions architect needs to assess a newly acquired company's portfolio of applications and databases. The solutions architect must create a business case to migrate the portfolio to AWS. The newly acquired company runs applications in an on-premises data center. The data center is not well documented. The solutions architect cannot immediately determine how many applications and databases exist. Traffic for the applications is variable. Some applications are batch processes that run at the end of each month.
The solutions architect must gain a better understanding of the portfolio before a migration to AWS can begin.
Which solution will meet these requirements?

A. Use AWS Server Migration Service (AWS SMS) and AWS Database Migration Service (AWS DMS) to evaluate migration. Use AWS Service Catalog to understand application and database dependencies.

B. Use AWS Application Migration Service. Run agents on the on-premises infrastructure. Manage the agents by using AWS Migration Hub. Use AWS Storage Gateway to assess local storage needs and database dependencies.

C. Use Migration Evaluator to generate a list of servers. Build a report for a business case. Use AWS Migration Hub to view the portfolio. Use AWS Application Discovery Service to gain an understanding of application dependencies.

D. Use AWS Control Tower in the destination account to generate an application portfolio. Use AWS Server Migration Service (AWS SMS) to generate deeper reports and a business case. Use a landing zone for core accounts and resources.

Correct Answer: C
Section:
Explanation:
The company should use Migration Evaluator to generate a list of servers and build a report for a business case. The company should use AWS Migration Hub to view the portfolio and use AWS Application Discovery Service to gain an understanding of application dependencies. This solution will meet the requirements because Migration Evaluator is a migration assessment service that helps create a data-driven business case for AWS cloud planning and migration.Migration Evaluator provides a clear baseline of what the company is running today and projects AWS costs based on measured on-premises provisioning and utilization1.Migration Evaluator can use an agentless collector to conduct broad-based discovery or securely upload exports from existing inventory tools2.Migration Evaluator integrates with AWS Migration Hub, which is a service that provides a single location to track the progress of application migrations across multiple AWS and partner solutions3.Migration Hub also supports AWS Application Discovery Service, which is a service that helps systems integrators quickly and reliably plan application migration projects by automatically identifying applications running in on-premises data centers, their associated dependencies, and their performance profile4.
https://aws.amazon.com/migration-evaluator/
https://docs.aws.amazon.com/migration-evaluator/latest/userguide/what-is.html
https://aws.amazon.com/migration-hub/
https://aws.amazon.com/application-discovery/
https://aws.amazon.com/server-migration-service/
https://aws.amazon.com/dms/
https://docs.aws.amazon.com/controltower/latest/userguide/what-is-control-tower.html
https://aws.amazon.com/application-migration-service/
https://aws.amazon.com/storagegateway/

QUESTION 114
A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.
Which solution will meet these requirements while providing the FASTEST storage performance?

A. Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluster. Configure the ReplicaSet to mount the file system. Direct the application to store files in the file system. Configure AWS Backup to back up and retain copies of the data for 1 year.
B. Create an Amazon Elastic Block Store (Amazon EBS) volume. Enable the EBS Multi-Attach feature. Configure the ReplicaSet to mount the EBS volume. Direct the application to store files in the EBS volume. Configure AWS Backup to back up and retain copies of the data for 1 year.
C. Create an Amazon S3 bucket. Configure the ReplicaSet to mount the S3 bucket. Direct the application to store files in the S3 bucket. Configure S3 Versioning to retain copies of the data. Configure an S3 Lifecycle policy to delete objects after 1 year.
D. Configure the ReplicaSet to use the storage available on each of the running application pods to store the files locally. Use a third-party tool to back up the EKS cluster for 1 year.

Correct Answer: A
Section:
Explanation:
In the past, EBS can be attached only to one ec2 instance but not anymore but there are limitations like - it works only on io1/io2 instance types and many others as described here.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html EFS has shareable storage In terms of performance, Amazon EFS is optimized for workloads that require high levels of aggregate throughput and IOPS, whereas EBS is optimized for low-latency, random access I/O operations. Amazon EFS is designed to scale throughput and capacity automatically as your storage needs grow, while EBS volumes can be resized on demand.

QUESTION 115
A company runs a customer service center that accepts calls and automatically sends all customers a managed, interactive, two-way experience survey by text message.
The applications that support the customer service center run on machines that the company hosts in an on-premises data center. The hardware that the company uses is old, and the company is experiencing downtime with the system. The company wants to migrate the system to AWS to improve reliability.
Which solution will meet these requirements with the LEAST ongoing operational overhead?

A. Use Amazon Connect to replace the old call center hardware. Use Amazon Pinpoint to send text message surveys to customers.
B. Use Amazon Connect to replace the old call center hardware. Use Amazon Simple Notification Service (Amazon SNS) to send text message surveys to customers.
C. Migrate the call center software to Amazon EC2 instances that are in an Auto Scaling group. Use the EC2 instances to send text message surveys to customers.
D. Use Amazon Pinpoint to replace the old call center hardware and to send text message surveys to customers.

Correct Answer: A
Section:
Explanation:
Amazon Connect is a cloud-based contact center service that allows you to set up a virtual call center for your business. It provides an easy-to-use interface for managing customer interactions through voice and chat. Amazon Connect integrates with other AWS services, such as Amazon S3 and Amazon Kinesis, to help you collect, store, and analyze customer data for insights into customer behavior and trends. On the other hand, Amazon Pinpoint is a marketing automation and analytics service that allows you to engage with your customers across different channels, such as email, SMS, push notifications, and voice. It helps you create personalized campaigns based on user behavior and enables you to track user engagement and retention. While both services allow you to communicate with your customers, they serve different purposes. Amazon Connect is focused on customer support and service, while Amazon Pinpoint is focused on marketing and engagement.

QUESTION 116
A company is building a call center by using Amazon Connect. The company's operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers.
Which solution will provide DR with the LOWEST RTO?

A. Create an AWS Lambda function to check the availability of the Amazon Connect instance and to send a notification to the operations team in case of unavailability. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes. After notification, instruct the operations team to use the AWS Management Console to provision a new Amazon Connect instance in a second Region. Deploy the contact flows, users, and claimed phone numbers by using an AWS CloudFormation template.
B. Provision a new Amazon Connect instance with all existing users in a second Region. Create an AWS Lambda function to check the availability of the Amazon Connect instance. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes. In the event of an issue, configure the Lambda function to deploy an AWS CloudFormation template that provisions contact flows and claimed numbers in the second Region.
C. Provision a new Amazon Connect instance with all existing contact flows and claimed phone numbers in a second Region. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance. Create an Amazon CloudWatch alarm for failed health checks. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions all users. Configure the alarm to invoke the Lambda function.
D. Provision a new Amazon Connect instance with all existing users and contact flows in a second Region. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance. Create an Amazon

CloudWatch alarm for failed health checks. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers. Configure the alarm to invoke the Lambda function.

Correct Answer: D
Section:
Explanation:
Option D provisions a new Amazon Connect instance with all existing users and contact flows in a second Region. It also sets up an Amazon Route 53 health check for the URL of the Amazon Connect instance, an Amazon CloudWatch alarm for failed health checks, and an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers. This option allows for the fastest recovery time because all the necessary components are already provisioned and ready to go in the second Region. In the event of a disaster, the failed health check will trigger the AWS Lambda function to deploy the CloudFormation template to provision the claimed phone numbers, which is the only missing component.

QUESTION 117
A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.
The company has the following DNS resolution requirements:
* On-premises systems should be able to resolve and connect to cloud.example.com.
* All VPCs should be able to resolve cloud.example.com.
There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway. Which architecture should the company use to meet these requirements with the HIGHEST performance?

A. Associate the private hosted zone to all the VPCs. Create a Route 53 inbound resolver in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.
B. Associate the private hosted zone to all the VPCs. Deploy an Amazon EC2 conditional forwarder in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the conditional forwarder.
C. Associate the private hosted zone to the shared services VPC. Create a Route 53 outbound resolver in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the outbound resolver.
D. Associate the private hosted zone to the shared services VPC. Create a Route 53 inbound resolver in the shared services VPC. Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

Correct Answer: A
Section:
Explanation:
Amazon Route 53 Resolver is a managed DNS resolver service from Route 53 that helps to create conditional forwarding rules to redirect query traffic1. By associating the private hosted zone to all the VPCs, the solutions architect can enable DNS resolution for cloud.example.com within the VPCs. By creating a Route 53 inbound resolver in the shared services VPC, the solutions architect can enable DNS resolution for cloud.example.com from on-premises systems. By attaching all VPCs to the transit gateway, the solutions architect can enable connectivity between the VPCs and the on-premises network through AWS Direct Connect. By creating forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver, the solutions architect can direct DNS queries for cloud.example.com to the Route 53 Resolver endpoint in AWS. This solution will provide the highest performance as it leverages Route 53 Resolver's optimized routing and caching capabilities.

QUESTION 118
A company is providing weather data over a REST-based API to several customers. The API is hosted by Amazon API Gateway and is integrated with different AWS Lambda functions for each API operation. The company uses Amazon Route 53 for DNS and has created a resource record of weather.example.com. The company stores data for the API in Amazon DynamoDB tables. The company needs a solution that will give the API the ability to fail over to a different AWS Region.
Which solution will meet these requirements?

A. Deploy a new set of Lambda functions in a new Region. Update the API Gateway API to use an edge-optimized API endpoint with Lambda functions from both Regions as targets. Convert the DynamoDB tables to global tables.
B. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a multivalue answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.
C. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a failover record. Enable target health monitoring. Convert the DynamoDB tables to global tables.
D. Deploy a new API Gateway API in a new Region. Change the Lambda functions to global functions. Change the Route 53 DNS record to a multivalue answer. Add both API Gateway APIs to the answer. Enable target health

monitoring. Convert the DynamoDB tables to global tables.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/apigateway/latest/developerguide/dns-failover.html

QUESTION 119
A solutions architect is designing a solution to process events. The solution must have the ability to scale in and out based on the number of events that the solution receives. If a processing error occurs, the event must move into a separate queue for review.
Which solution will meet these requirements?

A. Send event details to an Amazon Simple Notification Service (Amazon SNS) topic. Configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Add an on-failure destination to the function. Set an Amazon Simple Queue Service (Amazon SQS) queue as the target.
B. Publish events to an Amazon Simple Queue Service (Amazon SQS) queue. Create an Amazon EC2 Auto Scaling group. Configure the Auto Scaling group to scale in and out based on the ApproximateAgeOfOldestMessage metric of the queue. Configure the application to write failed messages to a dead-letter queue.
C. Write events to an Amazon DynamoDB table. Configure a DynamoDB stream for the table. Configure the stream to invoke an AWS Lambda function. Configure the Lambda function to process the events.
D. Publish events to an Amazon EventBridge event bus. Create and run an application on an Amazon EC2 instance with an Auto Scaling group that is behind an Application Load Balancer (ALB). Set the ALB as the event bus target. Configure the event bus to retry events. Write messages to a dead-letter queue if the application cannot process the messages.

Correct Answer: A
Section:
Explanation:
Amazon Simple Notification Service (Amazon SNS) is a fully managed pub/sub messaging service that enables users to send messages to multiple subscribers1. Users can send event details to an Amazon SNS topic and configure an AWS Lambda function as a subscriber to the SNS topic to process the events.Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources2. Users can add an on-failure destination to the function and set an Amazon Simple Queue Service (Amazon SQS) queue as the target.Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices, distributed systems, and serverless applications3. This way, if a processing error occurs, the event will move into the separate queue for review.
Option B is incorrect because publishing events to an Amazon SQS queue and creating an Amazon EC2 Auto Scaling group will not have the ability to scale in and out based on the number of events that the solution receives. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Auto Scaling is a feature that helps users maintain application availability and allows them to scale their EC2 capacity up or down automatically according to conditions they define. However, for this use case, using SQS and EC2 will not take advantage of the serverless capabilities of Lambda and SNS.
Option C is incorrect because writing events to an Amazon DynamoDB table and configuring a DynamoDB stream for the table will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB Streams is a feature that captures data modification events in DynamoDB tables. Users can configure the stream to invoke a Lambda function, but they cannot configure an on-failure destination for the function.
Option D is incorrect because publishing events to an Amazon EventBridge event bus and setting an Application Load Balancer (ALB) as the event bus target will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. An ALB is a load balancer that distributes incoming application traffic across multiple targets, such as EC2 instances, containers, IP addresses, Lambda functions, and virtual appliances. Users can configure EventBridge to retry events, but they cannot configure an on-failure destination for the ALB.

QUESTION 120
A company has loT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.
A solutions architect designs a solution in which the loT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically and are periodically encountering a RealProvisioned Throughput Exceeded error.
Which actions should the solution architect take to resolve this issue? (Select THREE.)

A. Reshard the stream to increase the number of shards s in the stream.
B. Use the Kinesis Producer Library KPL). Adjust the polling frequency.
C. Use consumers with the enhanced fan-out feature.
D. Reshard the stream to reduce the number of shards in the stream.

E. Use an error retry and exponential backoff mechanism in the consumer logic.
F. Configure the stream to use dynamic partitioning.

Correct Answer: A, C, E
Section:
Explanation:
https://repost.aws/knowledge-center/kinesis-readprovisionedthroughputexceeded
Follow Data Streams best practices
To mitigate ReadProvisionedThroughputExceeded exceptions, apply these best practices:
* Reshard your stream to increase the number of shards in the stream.
* Use consumers with enhanced fan-out. For more information about enhanced fan-out, see Developing custom consumers with dedicated throughput (enhanced fan-out).
* Use an error retry and exponential backoff mechanism in the consumer logic if ReadProvisionedThroughputExceeded exceptions are encountered. For consumer applications that use an AWS SDK, the requests are retried by default.

QUESTION 121
An education company is running a web application used by college students around the world. The application runs in an Amazon Elastic Container Service (Amazon ECS) cluster in an Auto Scaling group behind an Application Load Balancer (ALB). A system administrator detected a weekly spike in the number of failed logic attempts. Which overwhelm the application's authentication service. All the failed login attempts originate from about 500 different IP addresses that change each week. A solutions architect must prevent the failed login attempts from overwhelming the authentication service.
Which solution meets these requirements with the MOST operational efficiency?

A. Use AWS Firewall Manager to create a security group and security group policy to deny access from the IP addresses.
B. Create an AWS WAF web ACL with a rate-based rule, and set the rule action to Block. Connect the web ACL to the ALB.
C. Use AWS Firewall Manager to create a security group and security group policy to allow access only to specific CIDR ranges.
D. Create an AWS WAF web ACL with an IP set match rule, and set the rule action to Block. Connect the web ACL to the ALB.

Correct Answer: B
Section:

QUESTION 122
A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances.
To meet regulatory and business requirements, the company must make the following changes for data backups:
* Backups must be retained based on custom daily, weekly, and monthly requirements.
* Backups must be replicated to at least one other AWS Region immediately after capture.
* The backup solution must provide a single source of backup status across the AWS environment.
* The backup solution must send immediate notifications upon failure of any resource backup.
Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Select THREE.)

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

Correct Answer: A, B, D
Section:
Explanation:
Cross region with AWS Backup: https://docs.aws.amazon.com/aws-backup/latest/devguide/cross-region-backup.html

QUESTION 123
A company has many separate AWS accounts and uses no central billing or management. Each AWS account hosts services for different departments in the company. The company has a Microsoft Azure Active Directory that is deployed.
A solution architect needs to centralize billing and management of the company's AWS accounts. The company wants to start using identify federation instead of manual user management. The company also wants to use temporary credentials instead of long-lived access keys.
Which combination of steps will meet these requirements? (Select THREE)

A. Create a new AWS account to serve as a management account. Deploy an organization in AWS Organizations. Invite each existing AWS account to join the organization. Ensure that each account accepts the invitation.
B. Configure each AWS Account's email address to be aws+<account id>@example.com so that account management email messages and invoices are sent to the same place.
C. Deploy AWS IAM Identity Center (AWS Single Sign-On) in the management account. Connect IAM Identity Center to the Azure Active Directory. Configure IAM Identity Center for automatic synchronization of users and groups.
D. Deploy an AWS Managed Microsoft AD directory in the management account. Share the directory with all other accounts in the organization by using AWS Resource Access Manager (AWS RAM).
E. Create AWS IAM Identity Center (AWS Single Sign-On) permission sets. Attach the permission sets to the appropriate IAM Identity Center groups and AWS accounts.
F. Configure AWS Identity and Access Management (IAM) in each AWS account to use AWS Managed Microsoft AD for authentication and authorization.

Correct Answer: A, C, E
Section:

QUESTION 124
A company processes environment data. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.
The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time.
Which solution will meet these requirements?

A. Use Amazon Kinesis Data Firehouse to send the data to Amazon Redshift.
B. Use Amazon Kinesis Data streams to send the data to Amazon DynamoDB.
C. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to send the data to Amazon Aurora.
D. Use Amazon Kinesis Data firehouse to send the data to Amazon Keyspaces (for Apache Cassandra).

Correct Answer: B
Section:
Explanation:
Amazon Kinesis Data Streams is a service that enables real-time data ingestion and processing. Amazon DynamoDB is a NoSQL database that does not require fixed schemas for storage. By using Kinesis Data Streams and DynamoDB, the company can send the JSON data to a database that can handle schemaless data in real time.
Reference:
https://docs.aws.amazon.com/streams/latest/dev/introduction.html
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html

QUESTION 125
A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-through put. low-latency network connections between all to the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant.
Which solution will meet these requirements?

A. Launch five new EC2 instances into a cluster placement group. Ensure that the EC2 instance type supports enhanced networking.
B. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zone. Attach an extra elastic network interface to each EC2 instance.
C. Launch five new EC2 instances into a partition placement group. Ensure that the EC2 instance type supports enhanced networking.
D. Launch five new EC2 instances into a spread placement group Attach an extra elastic network interface to each EC2 instance.

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html#placement- groups-cluster

QUESTION 126
A company has migrated a legacy application to the AWS Cloud. The application runs on three Amazon EC2 instances that are spread across three Availability Zones. One EC2 instance is in each Availability Zone. The EC2 instances are running in three private subnets of the VPC and are set up as targets for an Application Load Balancer (ALB) that is associated with three public subnets.
The application needs to communicate with on-premises systems. Only traffic from IP addresses in the company's IP address range are allowed to access the on-premises systems. The company's security team is bringing only one IP address from its internal IP address range to the cloud. The company has added this IP address to the allow list for the company firewall. The company also has created an Elastic IP address for this IP address.
A solutions architect needs to create a solution that gives the application the ability to communicate with the on-premises systems. The solution also must be able to mitigate failures automatically.
Which solution will meet these requirements?

A. Deploy three NAT gateways, one in each public subnet. Assign the Elastic IP address to the NAT gateways. Turn on health checks for the NAT gateways. If a NAT gateway fails a health check, recreate the NAT gateway and assign the Elastic IP address to the new NAT gateway.
B. Replace the ALB with a Network Load Balancer (NLB). Assign the Elastic IP address to the NLB Turn on health checks for the NLB. In the case of a failed health check, redeploy the NLB in different subnets.
C. Deploy a single NAT gateway in a public subnet. Assign the Elastic IP address to the NAT gateway. Use Amazon CloudWatch with a custom metric to monitor the NAT gateway. If the NAT gateway is unhealthy, invoke an AWS Lambda function to create a new NAT gateway in a different subnet. Assign the Elastic IP address to the new NAT gateway.
D. Assign the Elastic IP address to the ALB. Create an Amazon Route 53 simple record with the Elastic IP address as the value. Create a Route 53 health check. In the case of a failed health check, recreate the ALB in different subnets.

Correct Answer: C
Section:
Explanation:
to connect out from the private subnet you need an NAT gateway and since only one Elastic IP whitelisted on firewall its one NATGateway at time and if AZ failure happens Lambda creates a new NATGATEWAY in a different AZ using the Same Elastic IP ,dont be tempted to select D since application that needs to connect is on a private subnet whose outbound connections use the NATGateway Elastic IP

QUESTION 127
A company is creating a centralized logging service running on Amazon EC2 that will receive and analyze logs from hundreds of AWS accounts. AWS PrivateLink is being used to provide connectivity between the client services and the logging service.
In each AWS account with a client, an interface endpoint has been created for the logging service and is available. The logging service running on EC2 instances with a Network Load Balancer (NLB) are deployed in different subnets. The clients are unable to submit logs using the VPC endpoint.
Which combination of steps should a solutions architect take to resolve this issue? (Select TWO.)

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

Correct Answer: A, C
Section:

QUESTION 128
A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.

Which solution will meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VPC. Create VPC peering connections that initiate from the central VPC to all other VPCs.
B. Create an AWS Direct Connect connection between the on-premises data center and AWS. Provision a transit VIF, and connect it to a Direct Connect gateway. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.
C. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VPC. Use a transit gateway with dynamic routing. Connect the transit gateway to all other VPCs.
D. Create an AWS Direct Connect connection between the on-premises data center and AWS Establish an AWS Site-to-Site VPN connection between all VPCs in each Region. Create VPC peering connections that initiate from the central VPC to all other VPCs.

Correct Answer: B
Section:
Explanation:
Transit GW + Direct Connect GW + Transit VIF + enabled SiteLink if two different DX locations https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-aws-direct-connect-sitelink/

QUESTION 129
A company uses a load balancer to distribute traffic to Amazon EC2 instances in a single Availability Zone. The company is concerned about security and wants a solutions architect to re-architect the solution to meet the following requirements:
* Inbound requests must be filtered for common vulnerability attacks.
* Rejected requests must be sent to a third-party auditing application.
* All resources should be highly available.
Which solution meets these requirements?

A. Configure a Multi-AZ Auto Scaling group using the application's AMI. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Use Amazon Inspector to monitor traffic to the ALB and EC2 instances. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application.
B. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Logs. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.
C. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as targets. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.
D. Configure a Multi-AZ Auto Scaling group using the application's AMI. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the WebACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

Correct Answer: D
Section:
Explanation:
https://docs.aws.amazon.com/waf/latest/developerguide/marketplace-managed-rule-groups.html

QUESTION 130
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 togs 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 SOL 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 131
A company has built a high performance computing (HPC) cluster in AWS tor 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 132
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

QUESTION 133
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 134
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 135
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 tor 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 Spool 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 136
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 137
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 learn 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 (he creation of resources in Regions that are outside the United States. Apply the SCP to the OU.

E. Create an 1AM role in the management account Attach a policy that includes permissions to view the Billing and Cost Management console. Allow the finance learn users to assume the role. Use AWS Cost Explorer and the Billing and Cost Management console to analyze cost.

F. Create an 1AM 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_manage_policies_scp.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
https://docs.aws.amazon.com/aws-cost-management/latest/userguide/what-is-costmanagement.html

QUESTION 138
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 seventy 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 lo 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 139
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 lo Amazon S3. During a disaster, recover the on-premises servers from the backup. During failback. run the on-premises servers on Amazon EC2 instances.
B. Create a set of AWS CloudFormation templates to create infrastructure. Replicate all data to Amazon Elastic File System (Amazon EFS) by using AWS DataSync. During a disaster, use AWS CodePipeline to deploy the templates to restore the on-premises servers. Fail back the data by using DataSync.
C. Create an AWS Cloud Development Kit (AWS CDK) pipeline to stand up a multi-site active-active environment on AWS. Replicate data into Amazon S3 by using the s3 sync command. During a disaster, swap DNS endpoints to point to AWS. Fail back the data by using the s3 sync command.
D. Use AWS Elastic Disaster Recovery to replicate the on-premises servers. Replicate data to an Amazon FSx for Windows File Server file system by using AWS DataSync. Mount the file system to AWS servers. During a disaster, fail over the on-premises servers to AWS. Fail back to new or existing servers by using Elastic Disaster Recovery.

Correct Answer: D
Section:

QUESTION 140
A delivery company is running a serverless solution in tneAWS Cloud The solution manages user data, delivery information and past purchase details The solution consists of several microservices The central user service stores sensitive data in an Amazon DynamoDB table Several of the other microservices store a copy of parts of the sensitive data in different storage services
The company needs the ability to delete user information upon request As soon as the central user service deletes a user every other microservice must also delete its copy of the data immediately
Which solution will meet these requirements?

A. Activate DynamoDB Streams on the DynamoDB table Create an AWS Lambda trigger for the DynamoDB stream that will post events about user deletion in an Amazon Simple Queue Service (Amazon SQS) queue Configure each microservice to poll the queue and delete the user from the DynamoDB table
B. Set up DynamoDB event notifications on the DynamoDB table Create an Amazon Simple Notification Service (Amazon SNS) topic as a target for the DynamoDB event notification Configure each microservice to subscribe to the SNS topic and to delete the user from the DynamoDB table
C. Configure the central user service to post an event on a custom Amazon EventBridge event bus when the company deletes a user Create an EventBndge rule for each microservice to match the user deletion event pattern and invoke logic in the microservice to delete the user from the DynamoDB table
D. Configure the central user service to post a message on an Amazon Simple Queue Service (Amazon SQS) queue when the company deletes a user Configure each microservice to create an event filter on the SQS queue and to delete the user from the DynamoDB table

Correct Answer: C
Section:
Explanation:
Set Up EventBridge Event Bus:
Step 1: Open the Amazon EventBridge console and create a custom event bus. This bus will be used to handle user deletion events.
Step 2: Name the event bus appropriately (e.g., user-deletion-bus).
Post Events on User Deletion:
Step 1: Modify the central user service to post an event to the custom EventBridge event bus whenever a user is deleted.
Step 2: Ensure the event includes relevant details such as the user ID and any other necessary metadata.
Create EventBridge Rules for Microservices:
Step 1: For each microservice that needs to delete user data, create a new rule in EventBridge that triggers on the user deletion event.
Step 2: Define the event pattern to match the user deletion event. This pattern should include the event details posted by the central user service.
Invoke Microservice Logic:
Step 1: Configure the EventBridge rule to invoke a target, such as an AWS Lambda function, which contains the logic to delete the user data from the microservice's data store.
Step 2: Each microservice should have its Lambda function or equivalent logic to handle the deletion of user data upon receiving the event.
Using Amazon EventBridge ensures a scalable, reliable, and decoupled approach to handle the deletion of user data across multiple microservices. This setup allows each microservice to independently process user deletion events without direct dependencies on other services.
Reference
AWS EventBridge Documentation

DynamoDB Streams and AWS Lambda Triggers
Implementing the Transactional Outbox Pattern with EventBridge Pipes (AWS Documentation) (Amazon Web Services, Inc.) (Amazon Web Services, Inc.) (AWS Documentation) (AWS Cloud Community).

QUESTION 141
To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network The security team mandates that no traffic accessing this data should traverse the public internet How should the solutions architect design a highly available solution that meets the requirements and is cost-effective'?

A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use the company WAN to send traffic over to the headquarters and then to the respective DX connection to access the data
B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region Use the company WAN to send traffic over a DX connection Use inter-region VPC peering to access the data in other AWS Regions
C. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region Use the company WAN to send traffic over a DX connection Use an AWS transit VPC solution to access data in other AWS Regions
D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region Use the company WAN to send traffic over a DX connection Use Direct Connect Gateway to access data in other AWS Regions.

Correct Answer: D
Section:
Explanation:
Establish AWS Direct Connect Connections:
Step 1: Set up two AWS Direct Connect (DX) connections from the company headquarters to a chosen AWS Region. This provides a redundant and high-availability setup to ensure continuous connectivity.
Step 2: Ensure that these DX connections terminate in a specific Direct Connect location associated with the chosen AWS Region.
Use Company WAN:
Step 1: Configure the company's global WAN to route traffic through the established Direct Connect connections.
Step 2: This setup ensures that all traffic between the company's headquarters and AWS does not traverse the public internet, maintaining compliance with security requirements.
Set Up Direct Connect Gateway:
Step 1: Create a Direct Connect Gateway in the AWS Management Console. This gateway allows you to connect your Direct Connect connections to multiple VPCs across different AWS Regions.
Step 2: Associate the Direct Connect Gateway with the VPCs in the various Regions where your critical data is stored. This enables access to data in multiple Regions through a single Direct Connect connection.
By using Direct Connect and Direct Connect Gateway, the company can achieve secure, reliable, and cost-effective access to data stored across multiple AWS Regions without using the public internet, ensuring compliance with industry regulations.
Reference
AWS Direct Connect Documentation
Building a Scalable and Secure Multi-VPC AWS Network Infrastructure (AWS Documentation) (AWS Documentation).

QUESTION 142
A company has developed an application that is running Windows Server on VMware vSphere VMs that the company hosts on premises The application data is stored in a proprietary format that must be read through the application The company manually provisioned the servers and the application
As part of its disaster recovery plan, the company wants the ability to host its application on AWS temporarily if the company's on-premises environment becomes unavailable The company wants the application to return to on-premises hosting after a disaster recovery event is complete The RPO is 5 minutes.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Configure AWS DataSync Replicate the data to Amazon Elastic Block Store (Amazon EBS) volumes When the on-premises environment is unavailable, use AWS Cloud Format ion templates to provision Amazon EC2 instances and attach the EBS volumes
B. Configure AWS Elastic Disaster Recovery Replicate the data to replication Amazon EC2 instances that are attached to Amazon Elastic Block Store (Amazon EBS) volumes When the on-premises environment is unavailable use Elastic Disaster Recovery to launch EC2 instances that use the replicated volumes
C. Provision an AWS Storage Gateway file gateway. Replicate the data to an Amazon S3 bucket When the on-premises environment is unavailable, use AWS Backup to restore the data to Amazon Elastic Block Store (Amazon EBS) volumes and launch Amazon EC2 instances from these EBS volumes
D. Provision an Amazon FSx for Windows File Server file system on AWS Replicate the data to the file system When the on-premises environment is unavailable, use AWS Cloud Format ion templates to provision Amazon EC2

instances and use AWS CloudFormation Init commands to mount the Amazon FSx file shares

Correct Answer: B
Section:
Explanation:
Set Up AWS Elastic Disaster Recovery:
Navigate to the AWS Elastic Disaster Recovery (DRS) console.
Configure the Elastic Disaster Recovery service to replicate your on-premises VMware vSphere VMs to Amazon EC2 instances. This involves installing the AWS Replication Agent on your VMs.
Configure Replication Settings:
Define the replication settings, including the Amazon EC2 instance type and the Amazon EBS volume configuration. Ensure that the replication frequency meets your Recovery Point Objective (RPO) of 5 minutes.
Monitor Data Replication:
Monitor the initial data replication process in the Elastic Disaster Recovery console. Once the initial sync is complete, the status should show as 'Healthy' indicating that the data replication is up-to-date and within the RPO requirements.
Disaster Recovery (Failover):
In the event of a disaster, initiate a failover from the Elastic Disaster Recovery console. This will launch the replicated Amazon EC2 instances using the Amazon EBS volumes with the latest data.
Failback Process:
Once the on-premises environment is restored, perform a failback operation to synchronize the data from AWS back to your on-premises VMware environment. Use the failback client provided by AWS Elastic Disaster Recovery to ensure data consistency and minimal downtime during the failback process.
Using AWS Elastic Disaster Recovery provides a low-overhead, automated solution for disaster recovery that ensures minimal data loss and meets the RPO requirement of 5 minutes (Amazon Web Services, Inc.) (AWS Documentation).

QUESTION 143
A company needs to improve the security of its web-based application on AWS. The application uses Amazon CloudFront with two custom origins. The first custom origin routes requests to an Amazon API Gateway HTTP API. The second custom origin routes traffic to an Application Load Balancer (ALB) The application integrates with an OpenID Connect (OIDC) identity provider (IdP) for user management.
A security audit shows that a JSON Web Token (JWT) authorizer provides access to the API The security audit also shows that the ALB accepts requests from unauthenticated users
A solutions architect must design a solution to ensure that all backend services respond to only authenticated users
Which solution will meet this requirement?

A. Configure the ALB to enforce authentication and authorization by integrating the ALB with the IdP Allow only authenticated users to access the backend services
B. Modify the CloudFront configuration to use signed URLs Implement a permissive signing policy that allows any request to access the backend services
C. Create an AWS WAF web ACL that filters out unauthenticated requests at the ALB level. Allow only authenticated traffic to reach the backend services.
D. Enable AWS CloudTrail to log all requests that come to the ALB Create an AWS Lambda function to analyze the togs and block any requests that come from unauthenticated users.

Correct Answer: A
Section:
Explanation:
Integrate ALB with OIDC IdP:
In the AWS Management Console, navigate to the Application Load Balancer (ALB) settings.
Configure the ALB to use the OpenID Connect (OIDC) IdP for authentication. This ensures that all requests routed through the ALB are authenticated using the IdP.
Set Up Authentication Rules:
Create a listener rule on the ALB that requires authentication. This rule will forward requests to the IdP for user authentication before allowing access to the backend services.
Restrict Unauthenticated Access:
Ensure the ALB only forwards requests to backend services if the user is authenticated. Unauthenticated requests should be blocked or redirected to the IdP for authentication.
Update CloudFront Configuration:
Modify the CloudFront distribution to forward authenticated requests to the ALB. Ensure that the ALB and API Gateway accept only requests coming through the CloudFront distribution to enforce consistent authentication and security.
By enforcing authentication at the ALB level, you ensure that all backend services are accessed only by authenticated users, enhancing the overall security of the web application

QUESTION 144

A company is running its solution on AWS in a manually created VPC. The company is using AWS CloudFormation to provision other parts of the infrastructure According to a new requirement the company must manage all infrastructure in an automatic way
What should the comp any do to meet this new requirement with the LEAST effort?

A. Create a new AWS Cloud Development Kit (AWS CDK) stack that strictly provisions the existing VPC resources and configuration Use AWS CDK to import the VPC into the stack and to manage the VPC
B. Create a CloudFormation stack set that creates the VPC Use the stack set to import the VPC into the stack
C. Create a new CloudFormation template that strictly provisions the existing VPC resources and configuration From the CloudFormation console, create a new stack by importing the existing resources
D. Create a new CloudFormation template that creates the VPC Use the AWS Serverless Application Model (AWS SAM) CLI to import the VPC

Correct Answer: C
Section:
Explanation:
Creating the Template:
Start by creating a CloudFormation template that includes all the VPC resources. This template should accurately reflect the current state and configuration of the VPC.
Using the CloudFormation Console:
Open the AWS Management Console and navigate to CloudFormation.
Choose 'Create stack' and then select 'With existing resources (import resources)'.
Specifying the Template:
Upload the previously created template or specify the Amazon S3 URL where the template is stored.
Identifying the Resources:
On the 'Identify resources' page, provide the identifiers for each VPC resource you wish to import. For example, for an AWS::EC2::VPC resource, use the VPC ID as the identifier.
Creating the Stack:
Complete the stack creation process by providing stack details and reviewing the changes. This will create a change set that includes the import operation.
Executing the Change Set:
Execute the change set to import the resources into the CloudFormation stack, making them managed by CloudFormation.
Verification and Drift Detection:
After the import is complete, run drift detection to ensure the actual configuration matches the template configuration.
This approach allows the company to manage their VPC and other resources via CloudFormation without the need to recreate resources, ensuring a smooth transition to automated infrastructure management.
Reference
Creating a stack from existing resources - AWS CloudFormation (AWS Documentation).
Generating templates for existing resources - AWS CloudFormation (AWS Documentation).
Bringing existing resources into CloudFormation management (AWS Documentation).

QUESTION 145
A company runs a software-as-a-service <SaaS) application on AWS The application consists of AWS Lambda functions and an Amazon RDS for MySQL Multi-AZ database During market events the application has a much higher workload than normal Users notice slow response times during the peak periods because of many database connections The company needs to improve the scalable performance and availability of the database Which solution meets these requirements'?

A. Create an Amazon CloudWatch alarm action that triggers a Lambda function to add an Amazon RDS for MySQL read replica when resource utilization hits a threshold
B. Migrate the database to Amazon Aurora, and add a read replica Add a database connection pool outside of the Lambda handler function
C. Migrate the database to Amazon Aurora and add a read replica Use Amazon Route 53 weighted records
D. Migrate the database to Amazon Aurora and add an Aurora Replica Configure Amazon RDS Proxy to manage database connection pools

Correct Answer: D
Section:
Explanation:
Migrate to Amazon Aurora:
Amazon Aurora is a MySQL-compatible, high-performance database designed to provide higher throughput than standard MySQL. Migrating the database to Aurora will enhance the performance and scalability of the database, especially under heavy workloads.

Add Aurora Replica:
Aurora Replicas provide read scalability and improve availability. Adding an Aurora Replica allows read operations to be distributed, thereby reducing the load on the primary instance and improving response times during peak periods.
Configure Amazon RDS Proxy:
Amazon RDS Proxy acts as an intermediary between the application and the Aurora database, managing connection pools efficiently. RDS Proxy reduces the overhead of opening and closing database connections, thus maintaining fewer active connections to the database and handling surges in database connections from the Lambda functions more effectively.
This configuration reduces the database's resource usage and improves its ability to handle high volumes of concurrent connections.
Reference
AWS Database Blog on RDS Proxy (Amazon Web Services, Inc.).
AWS Compute Blog on Using RDS Proxy with Lambda (Amazon Web Services, Inc.).

QUESTION 146
A company provides a centralized Amazon EC2 application hosted in a single shared VPC The centralized application must be accessible from client applications running in the VPCs of other business units The centralized application front end is configured with a Network Load Balancer (NLB) for scalability
Up to 10 business unit VPCs will need to be connected to the shared VPC Some ot the business unit VPC CIDR blocks overlap with the shared VPC and some overlap with each other Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only
Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

A. Create an AWS Transit Gateway Attach the shared VPC and the authorized business unit VPCs to the transit gateway Create a single transit gateway route table and associate it with all of the attached VPCs Allow automatic propagation of routes from the attachments into the route table Configure VPC routing tables to send traffic to the transit gateway
B. Create a VPC endpoint service using the centralized application NLB and enable the option to require endpoint acceptance Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service. Accept authorized endpoint requests from the endpoint service console.
C. Create a VPC peering connection from each business unit VPC to the shared VPC Accept the VPC peering connections from the shared VPC console Configure VPC routing tables to send traffic to the VPC peering connection
D. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPCs Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC Configure VPC routing tables to send traffic to the VPN connection

Correct Answer: B
Section:
Explanation:
Create VPC Endpoint Service:
In the shared VPC, create a VPC endpoint service using the Network Load Balancer (NLB) that fronts the centralized application.
Enable the option to require endpoint acceptance to control which business unit VPCs can connect to the service.
Set Up VPC Endpoints in Business Unit VPCs:
In each business unit VPC, create a VPC endpoint that points to the VPC endpoint service created in the shared VPC.
Use the service name of the endpoint service created in the shared VPC for configuration.
Accept Endpoint Requests:
From the VPC endpoint service console in the shared VPC, review and accept endpoint connection requests from authorized business unit VPCs. This ensures that only authorized VPCs can access the centralized application.
Configure Routing:
Update the route tables in each business unit VPC to direct traffic destined for the centralized application through the VPC endpoint.
This solution ensures secure, private connectivity between the business unit VPCs and the shared VPC, even if there are overlapping CIDR blocks. It leverages AWS PrivateLink and VPC endpoints to provide scalable and controlled access (AWS Documentation) (Amazon Web Services, Inc.).

QUESTION 147
A company is running a large containerized workload in the AWS Cloud. The workload consists of approximately 100 different services. The company uses Amazon Elastic Container Service (Amazon ECS) to orchestrate the workload.
Recently, the company's development team started using AWS Fargate instead of Amazon EC2 instances in the ECS cluster. In the past, the workload has come close to running the maximum number of EC2 instances that are available in the account.
The company is worried that the workload could reach the maximum number of ECS tasks that are allowed. A solutions architect must implement a solution that will notify the development team when Fargate reaches 80%

of the maximum number of tasks.
What should the solutions architect do to meet this requirement?

A. Use Amazon CloudWatch to monitor the Sample Count statistic for each service in the ECS cluster. Set an alarm for when the math expression sample count/SERVICE_QUOTA(service)'100 is greater than 80. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
B. Use Amazon CloudWatch to monitor service quotas that are published under the AWS/Usage metric namespace. Set an alarm for when the math expression metric/SERVICE_QUOTA(metric)*100 is greater than 80. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
C. Create an AWS Lambda function to poll detailed metrics from the ECS cluster. When the number of running Fargate tasks is greater than 80, invoke Amazon Simple Email Service (Amazon SES) to notify the development team.
D. Create an AWS Config rule to evaluate whether the Fargate SERVICE_OUOTA is greater than 80. Use Amazon Simple Email Service (Amazon SES) to notify the development team when the AWS Config rule is not compliant.

Correct Answer: B
Section:

QUESTION 148
A company has a new requirement to store all database backups in an isolated AWS account. The company is using AWS Organizations and has created a central write-once, read-many (WORM) account for the backups.
The company has 40 Amazon RDS tor MySQL databases in its production account. The databases are encrypted with the default RDS AWS Key Management Service (AWS KMS) key. RDS automated backups of the databases occur daily and have a retention period of 30 days.
Which solution will successfully copy the database backups to the central account?

A. Enable Organizations trusted access and backup policies for AWS Backup. Configure the central account as the delegated administrator for AWS Backup. Create 1AM policies and backup policies. Enable cross-account management. Create a backup vault in the central account. Create a KMS key for the backup vault and share the key with the production account. In the production account, restore the databases from a snapshot and apply the shared KMS key to the new DB instances. Create a backup plan in the central account to back up the databases to The backup vault.
B. Enable Organizations trusted access and backup policies for AWS Backup. Configure the central account as the delegated administrator for AWS Backup. Create 1AM policies and backup policies. Enable cross-account management. In the production account, share the default RDS KMS key with the central account. Create a backup vault in the central account. Apply the shared default RDS KMS key to the backup vault. Create a backup plan in the central account to back up the databases to the backup vault.
C. Create an Amazon EventBridge rule to invoke an AWS Lambda function every day. Program the Lambda function to decrypt the snapshots and to initiate a copy request of all unencrypted snapshots to the central account. After the copy job is complete, create a new KMS key. Use the new KMS key to encrypt the database snapshots in the central account.
D. Create an Amazon EventBridge rule to invoke an AWS Lambda function every day. In the production account, share the default RDS KMS key with the central account. Program the Lambda function to decrypt the snapshots and to initiate a copy request of all unencrypted snapshots to the central account. After the copy job is complete, encrypt the database snapshots with the shared default RDS KMS key in the central account.

Correct Answer: B
Section:

QUESTION 149
A company has a single AWS account. The company runs workloads on Amazon EC2 instances in multiple VPCs in one AWS Region. The company also runs workloads in an on-premises data center that connects to the company's AWS account by using AWS Direct Connect.
The company needs all EC2 instances in the VPCs to resolve DNS queries tor the internal.company.com domain to the authoritative DNS server that is located in the on-premises data center. The solution must use private communication between the VPCs and the on-premises network. All route tables, network ACLs. and security groups are configured correctly between AWS and the on-premises data center.
Which combination of actions will meet these requirements? (Select THREE)

A. Create an Amazon Route 53 inbound endpoint in all the workload VPCs.
B. Create an Amazon Route 53 outbound endpoint in one of the workload VPCs.
C. Create an Amazon Route 53 Resolver rule with the Forward type configured to forward queries for internal.company.com for the on-premises DNS server.
D. Create an Amazon Route 53 Resolver rule with the System type configured to forward queries for internal.company.com to the on-premises DNS server.
E. Associate the Amazon Route 53 Resolver rule with all the workload VPCs.
F. Associate the Amazon Route 53 Resolver rule with the workload VPC with the new Route 53 endpoint.

Correct Answer: C, E
Section:

QUESTION 150
A company is running several applications in the AWS Cloud. The applications are specific to separate business units in the company. The company is running the components of the applications in several AWS accounts that are in an organization in AWS Organizations.
Every cloud resource in the company's organization has a tag that is named BusinessUnit. Every lag already has the appropriate value of the business unit name.
The company needs to allocate its cloud costs to different business units. The company also needs to visualize the cloud costs for each business unit.
Which solution will meet these requirements?

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

Correct Answer: D
Section:

QUESTION 151
A company wants to create a single Amazon S3 bucket for its data scientists to store work-related documents. The company uses AWS 1AM Identity Center to authenticate all users. A group for the data scientists was created.
The company wants to give the data scientists access lo only their own work. The company also wants to create monthly reports that show which documents each user accessed.
Which combination of steps will meet these requirements? (Select TWO)

A. Create a custom 1AM Identity Center permission set to grant the data scientists access to an S3 bucket prefix that matches their username tag. Use a policy to limit access to paths with the S{aws:PrincipalTag/userName}/* condition.
B. Create an 1AM Identity Center role for the data scientists group that has Amazon S3 read access and write access. Add an S3 bucket policy that allows access to the 1AM Identity Center role.
C. Configure AWS CloudTrail to log S3 data events and deliver the logs to an S3 bucket. Use Amazon Athena to run queries on the CloudTrail logs in Amazon S3 and generate reports.
D. Configure AWS CloudTrail to log S3 management events to CloudWatch. Use Amazon Athena's CloudWatch connector to query the logs and generate reports.
E. Enable S3 access logging to EMR File System (EMRFS). Use Amazon S3 Select to query logs and generate reports.

Correct Answer: A, C
Section:

QUESTION 152
A company hosts a public software as a service (SaaS) application on Amazon EC2 instances that run Linux. The EC2 instances are in multiple Availability Zones behind an Application Load Balancer. The application uses an Amazon RDS Multi-AZ database to store application data, including user sessions.
The company needs to minimize the latency that is involved in storing and accessing the user sessions.
Which solution will meet this requirement?

A. Create an Amazon S3 bucket. Store the user sessions in S3 Standard storage.
B. Create an Amazon FSx for Windows File Server volume. Mount the volume on each EC2 instance. Store the user sessions in the volume.
C. Create a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume. Enable Multi-Attach. Attach the volume to each EC2 instance. Store the user sessions in the volume.
D. Create an Amazon ElastiCache for Redis cluster. Store the user sessions in the Redis cache.

Correct Answer: D
Section:

QUESTION 153
A company is migrating a monolithic on-premises .NET Framework production application to AWS. Application demand will grow exponentially in the next 6 months. The company must ensure that the application can scale appropriately.
The application currently connects to a Microsoft SQL Server transactional database. The company has well-documented source code for the application. Some business logic is contained within stored procedures.
A solutions architect must recommend a solution to redesign the application to meet the growth in demand.
Which solution will meet this requirement MOST cost-effectively?

A. Use Amazon API Gateway APIs and Amazon EC2 Spot Instances to rehost the application with a scalable microservices architecture. Deploy the EC2 instances in a cluster placement group. Configure EC2 Auto Scaling. Store the data and stored procedures in Amazon RDS for SQL Server.
B. Use AWS Application Migration Service to migrate the application to AWS Elastic Beanstalk. Deploy Elastic Beanstalk packages to configure and deploy the application as microservices. Deploy Elastic Beanstalk across multiple Availability Zones and configure auto scaling. Store the data and stored procedures in Amazon RDS for MySQL.
C. Migrate the applications by using AWS App2Container. Use AWS Fargate in multiple AWS Regions to host the containers. Use Amazon API Gateway APIs and AWS Lambda functions to call the containers. Store the data and stored procedures in Amazon DynamoDB Accelerator (DAX).
D. Use Amazon API Gateway APIs and AWS Lambda functions to decouple the application into microservices. Use the AWS Schema Conversion Tool (AWS SCT) to review and modify the stored procedures. Store the data in Amazon Aurora Serverless v2.

Correct Answer: D
Section:

QUESTION 154
A company is migrating its blog platform to AWS. The company's on-premises servers connect to AWS through an AWS Site-to-Site VPN connection. The blog content is updated several times a day by multiple authors and is served from a file share on a network-attached storage (NAS) server.
The company needs to migrate the blog platform without delaying the content updates. The company has deployed Amazon EC2 instances across multiple Availability Zones to run the blog platform behind an Application Load Balancer. The company also needs to move 200 TB of archival data from its on-premises servers to Amazon S3 as soon as possible.
Which combination of steps will meet these requirements? (Select TWO)

A. Create a weekly cron job in Amazon EventBridge. Use the cron job to invoke an AWS Lambda function to update the EC2 instances from the NAS server.
B. Configure an Amazon Elastic Block Store (Amazon EBS) Multi-Attach volume for the EC2 instances to share for content access. Write code to synchronize the EBS volume with the NAS server weekly.
C. Mount an Amazon Elastic File System (Amazon EFS) file system to the on-premises servers to act as the NAS server. Copy the blog data to the EPS file system. Mount the EFS file system to the EC2 instances to serve the content.
D. Order an AWS Snowball Edge Storage Optimized device. Copy the static data artifacts to the device. Ship the device to AWS.
E. Order an AWS Snowcone SSD device. Copy the static data artifacts to the device. Ship the device to AWS.

Correct Answer: C, D
Section:

QUESTION 155
A company is designing its network configuration in the AWS Cloud. The company uses AWS Organizations to manage a multi-account setup. The company has three OUs. Each OU contains more than 100 AWS accounts. Each account has a single VPC. and all the VPCs in each OU are in the same AWS Region.
The CIDR ranges for all the AWS accounts do not overlap. The company needs to implement a solution in which VPCs in the same OU can communicate with each other but cannot communicate with VPCs in other OUs.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS CloudFormation stack set that establishes VPC peering between accounts in each OU. Provision the stack set in each OU.
B. In each OU. create a dedicated networking account that has a single VPC. Share this VPC with all the other accounts in the OU by using AWS Resource Access Manager (AWS RAM). Create a VPC peering connection between the networking account and each account in the OU.

C. Provision a transit gateway in an account in each OU. Share the transit gateway across the organization by using AWS Resource Access Manager {AWS RAM). Create transit gateway VPC attachments for each VPC.
D. In each OU. create a dedicated networking account that has a single VPC. Establish a VPN connection between the networking account and the other accounts in the OU. Use third-party routing software to route transitive traffic between the VPCs.

Correct Answer: C
Section:

QUESTION 156
A company is using AWS Cloud Formation as its deployment tool for all applications. It stages all application binaries and templates within Amazon S3 buckets with versioning enabled. Developers have access to an Amazon EC2 instance that hosts the integrated development environment (IDE). The developers download the application binaries from Amazon S3 to the EC2 instance, make changes, and upload the binaries to an S3 bucket after running the unit tests locally. The developers want to improve the existing deployment mechanism and implement CI/CD using AWS CodePipeline.
The developers have the following requirements:
* Use AWS CodeCommit for source control.
* Automate unit testing and security scanning.
* Alert the developers when unit tests fail.
* Turn application features on and off, and customize deployment dynamically as part of CI/CD. Have the lead developer provide approval before deploying an application.
Which solution will meet these requirements?

A. Use AWS CodeBuild to run unit tests and security scans. Use an Amazon EventBridge rule to send Amazon SNS alerts to the developers when unit tests fail. Write AWS Cloud Development Kit (AWS CDK) constructs for different solution features, and use a manifest file to turn features on and off in the AWS CDK application. Use a manual approval stage in the pipeline to allow the lead developer to approve applications.
B. Use AWS Lambda to run unit tests and security scans. Use Lambda in a subsequent stage in the pipeline to send Amazon SNS alerts to the developers when unit tests fail. Write AWS Amplify plugins for different solution features and utilize user prompts to turn features on and off. Use Amazon SES in the pipeline to allow the lead developer to approve applications.
C. Use Jenkins to run unit tests and security scans. Use an Amazon EventBridge rule in the pipeline to send Amazon SES alerts to the developers when unit tests fail. Use AWS CloudFormation nested stacks for different solution features and parameters to turn features on and off. Use AWS Lambda in the pipeline to allow the lead developer to approve applications.
D. Use AWS CodeDeploy to run unit tests and security scans. Use an Amazon CloudWatch alarm in the pipeline to send Amazon SNS alerts to the developers when unit tests fail. Use Docker images for different solution features and the AWS CLI to turn features on and off. Use a manual approval stage in the pipeline to allow the lead developer to approve applications.

Correct Answer: A
Section:

QUESTION 157
A company has an application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The application is in an AWS account that has AWS CloudTrail enabled. The company restricts access to the application by adding the IP addresses of end users to a security group that is associated with the ALB.
The company is developing an AWS Lambda function to determine if the allowed IP addresses have accessed the application recently. If an allowed IP address has not accessed the application in the last 90 days, the Lambda function will remove the IP address from the security group.
The company needs to implement the functionality for the Lambda function to check the IP addresses.
Which combination of steps will provide this functionality MOST cost-effectively? (Select TWO)

A. For the VPC that contains (he ALB, configure VPC flow logs to be sent to a log group in Amazon CloudWatch Logs.
B. Enable access logging on the ALB. Create an Amazon Athena table to query the ALB access logs.
C. Program the Lambda function to check when each allowed IP address from the security group last appeared in the VPC flow logs.
D. Program the Lambda function to check when each allowed IP address from the security group last appeared in the ALB access logs
E. Program the Lambda function to check when each allowed IP address from the security group last appeared in the CloudTrail logs.

Correct Answer: B, D
Section:

QUESTION 158

A company has separate AWS accounts for each of its departments. The accounts are in OUs that are in an organization in AWS Organizations. The IT department manages a private certificate authority (CA) by using AWS Private Certificate Authority in its account.
The company needs a solution to allow developer teams in the other departmental accounts to access the private CA to issue certificates for their applications. The solution must maintain appropriate security boundaries between accounts.
Which solution will meet these requirements?

A. Create an AWS Lambda function in the IT account. Program the Lambda function to use theAWS Private CA API to export and import a private CA certificate to each department account. Use Amazon EventBridge to invoke the Lambda function on a schedule.
B. Create an 1AM identity-based policy that allows cross-account access to AWS Private CA. In the IT account, attach this policy to the private CA. Grant access to AWS Private CA by using the AWS Private CA API.
C. In the organization's management account, create an AWS CloudFormation stack to set up a resource-based delegation policy. Update the policy to allow the organizations: EnableAWSServiceAccess action. Add the ARN of the private CA from the IT account as the principal in the policy statement.
D. Use AWS Resource Access Manager (AWS RAM) in the IT account to enable sharing in the organization. Create a resource share. Add the private CA resource to the resource share. Grant the department OUs access to the shared CA.

Correct Answer: C
Section:

QUESTION 159
A company hosts its primary API on AWS by using an Amazon API Gateway API and AWS Lambda functions that contain the logic tor the API methods. The company's internal applications use the API tor core functionality and business logic. The company's customers use the API to access data from their accounts. Several customers also have access to a legacy API that is running on a single standalone Amazon EC2 instance.
The company wants to increase the security for these APIs to better prevent denial of service (DoS) attacks, check for vulnerabilities, and guard against common exploits.
What should a solutions architect do to meet these requirements?

A. Use AWS WAF to protect both APIs. Configure Amazon Inspector to analyze the legacy API. Configure Amazon GuardDuty to monitor (or malicious attempts to access the APIs.
B. Use AWS WAF to protect the API Gateway API. Configure Amazon Inspector to analyze both APIs. Configure Amazon GuardDuty to block malicious attempts to access the APIs.
C. Use AWS WAF to protect the API Gateway API. Configure Amazon Inspector to analyze the legacy API. Configure Amazon GuardDuty to monitor for malicious attempts to access the APIs.
D. Use AWS WAF to protect the API Gateway API. Configure Amazon Inspector to protect the legacy API. Configure Amazon GuardDuty to block malicious attempts to access the APIs.

Correct Answer: B
Section: