# Amazon

## Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

**NEW QUESTION 1**
- (Topic 4)
A company is creating an application that runs on containers in a VPC. The application stores and accesses data in an Amazon S3 bucket During the development phase, the application will store and access 1 TB of data in Amazon S3 each day. The company wants to minimize costs and wants to prevent traffic from traversing the internet whenever possible.
Which solution will meet these requirements?

A. Enable S3 Intelligent-Tiering for the S3 bucket.
B. Enable S3 Transfer Acceleration for the S3 bucket.
C. Create a gateway VPC endpoint for Amazon S3. Associate this endpoint with all route tables in the VPC.
D. Create an interface endpoint for Amazon S3 in the VP
E. Associate this endpoint with all route tables in the VPC.

**Answer:** C

**Explanation:**
A gateway VPC endpoint for Amazon S3 enables private connections between the VPC and Amazon S3 that do not require an internet gateway or NAT device. This minimizes costs and prevents traffic from traversing the internet. A gateway VPC endpoint uses a prefix list as the route target in a VPC route table to route traffic privately to Amazon S31. Associating the endpoint with all route tables in the VPC ensures that all subnets can access Amazon S3 through the endpoint.
Option A is incorrect because S3 Intelligent-Tiering is a storage class that optimizes storage costs by automatically moving objects between two access tiers based on changing access patterns. It does not affect the network traffic between the VPC and Amazon S32.
Option B is incorrect because S3 Transfer Acceleration is a feature that enables fast, easy, and secure transfers of files over long distances between clients and an S3 bucket. It does not prevent traffic from traversing the internet3.
Option D is incorrect because an interface VPC endpoint for Amazon S3 is powered by AWS PrivateLink, which requires an elastic network interface (ENI) with a private IP address in each subnet. This adds complexity and cost to the solution. Moreover, an interface VPC endpoint does not support cross-Region access to Amazon S3. Reference URL: 1: https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html 2: https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc- dynamic-data-access 3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html : https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for- amazon-s3/

**NEW QUESTION 2**
- (Topic 4)
A company has two VPCs named Management and Production. The Management VPC uses VPNs through a customer gateway to connect to a single device in the data center. The Production VPC uses a virtual private gateway AWS Direct Connect connections. The Management and Production VPCs both use a single VPC peering connection to allow communication between the
What should a solutions architect do to mitigate any single point of failure in this architecture?

A. Add a set of VPNs between the Management and Production VPCs.
B. Add a second virtual private gateway and attach it to the Management VPC.
C. Add a second set of VPNs to the Management VPC from a second customer gateway device.
D. Add a second VPC peering connection between the Management VPC and the Production VPC.

**Answer:** C

**Explanation:**
This answer is correct because it provides redundancy for the VPN connection between the Management VPC and the data center. If one customer gateway device or one VPN tunnel becomes unavailable, the traffic can still flow over the second customer gateway device and the second VPN tunnel. This way, the single point of failure in the VPN connection is mitigated.
References:
? https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-redundant-connection.html
? https://www.trendmicro.com/cloudoneconformity/knowledge-base/aws/VPC/vpn-tunnel-redundancy.html

**NEW QUESTION 3**
- (Topic 4)
A company wants to use an AWS CloudFormatlon stack for its application in a test environment. The company stores the CloudFormation template in an Amazon S3 bucket that blocks public access. The company wants to grant CloudFormation access to the template in the S3 bucket based on specific user requests to create the test environment The solution must follow security best practices.
Which solution will meet these requirements?

A. Create a gateway VPC endpoint for Amazon S3. Configure the CloudFormation stack to use the S3 object URL
B. Create an Amazon API Gateway REST API that has the S3 bucket as the targe
C. Configure the CloudFormat10n stack to use the API Gateway URL _
D. Create a presigned URL for the template object_ Configure the CloudFormation stack to use the presigned URL.
E. Allow public access to the template object in the S3 bucke
F. Block the public access after the test environment is created

**Answer:** C

**Explanation:**
it allows CloudFormation to access the template in the S3 bucket without granting public access or creating additional resources. A presigned URL is a URL that is signed with the access key of an IAM user or role that has permission to access the object. The presigned URL can be used by anyone who receives it, but it expires after a specified time. By creating a presigned URL for the template object and configuring the CloudFormation stack to use it, the company can grant CloudFormation access to the template based on specific user requests and follow security best practices. References:
? Using Amazon S3 Presigned URLs
? Using Amazon S3 Buckets

**NEW QUESTION 4**
- (Topic 4)

A company is moving its data and applications to AWS during a multiyear migration project. The company wants to securely access data on Amazon S3 from the company's AWS Region and from the company's on-premises location. The data must not traverse the internet. The company has established an AWS Direct Connect connection between its Region and its on-premises location
Which solution will meet these requirements?

A. Create gateway endpoints for Amazon S3. Use the gateway endpoints to securely access the data from the Region and the on-premises location.
B. Create a gateway in AWS Transit Gateway to access Amazon S3 securely from the Region and the on-premises location.
C. Create interface endpoints for Amazon S3_ Use the interface endpoints to securely access the data from the Region and the on-premises location.
D. Use an AWS Key Management Service (AWS KMS) key to access the data securely from the Region and the on-premises location.

**Answer:** B

**Explanation:**
A gateway endpoint is a gateway that is a target for a specified route in your route table, used for traffic destined to a supported AWS service1. Amazon S3 does not support gateway endpoints, only interface endpoints2. Therefore, option A is incorrect.
An interface endpoint is an elastic network interface with a private IP address that serves as an entry point for traffic destined to a supported service1. An interface endpoint can provide secure access to Amazon S3 from within the Region, but not from the on-premises location. Therefore, option C is incorrect.
AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys to protect your data3. AWS KMS does not provide a way to access data on Amazon S3 without traversing the internet. Therefore, option D is incorrect. AWS Transit Gateway is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. You can create a gateway in AWS Transit Gateway to access Amazon S3 securely from both the Region and the on-premises location using AWS Direct Connect. Therefore, option B is correct.

**NEW QUESTION 5**
- (Topic 4)
A company wants to analyze and generate reports to track the usage of its mobile app. The app is popular and has a global user base The company uses a custom report building program to analyze application usage.
The program generates multiple reports during the last week of each month. The program takes less than 10 minutes to produce each report. The company rarely uses the program to generate reports outside of the last week of each month. The company wants to generate reports in the least amount of time when the reports are requested.
Which solution will meet these requirements MOST cost-effectively?

A. Run the program by using Amazon EC2 On-Demand Instance
B. Create an Amazon EventBridge rule to start the EC2 instances when reports are requeste
C. Run the EC2 instances continuously during the last week of each month.
D. Run the program in AWS Lambd
E. Create an Amazon EventBridge rule to run a Lambda function when reports are requested.
F. Run the program in Amazon Elastic Container Service (Amazon ECS). Schedule Amazon ECS to run the program when reports are requested.
G. Run the program by using Amazon EC2 Spot Instance
H. Create an Amazon EventBridge rule to start the EC2 instances when reports are requeste
I. Run the EC2 instances continuously during the last week of each month.

**Answer:** B

**Explanation:**
This solution meets the requirements most cost-effectively because it leverages the serverless and event-driven capabilities of AWS Lambda and Amazon EventBridge. AWS Lambda allows you to run code without provisioning or managing servers, and you pay only for the compute time you consume. Amazon EventBridge is a serverless event bus service that lets you connect your applications with data from various sources and routes that data to targets such as AWS Lambda. By using Amazon EventBridge, you can create a rule that triggers a Lambda function to run the program when reports are requested, and you can also schedule the rule to run during the last week of each month. This way, you can generate reports in the least amount of time and pay only for the resources you use.
References:
? AWS Lambda
? Amazon EventBridge

**NEW QUESTION 6**
- (Topic 4)
A company containerized a Windows job that runs on .NET 6 Framework under a Windows container. The company wants to run this job in the AWS Cloud. The job runs every 10 minutes. The job's runtime varies between 1 minute and 3 minutes.
Which solution will meet these requirements MOST cost-effectively?

A. Create an AWS Lambda function based on the container image of the jo
B. Configure Amazon EventBridge to invoke the function every 10 minutes.
C. Use AWS Batch to create a job that uses AWS Fargate resource
D. Configure the job scheduling to run every 10 minutes.
E. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the jo
F. Create a scheduled task based on the container image of the job to run every 10 minutes.
G. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the jo
H. Create a standalone task based on the container image of the jo
I. Use Windows task scheduler to run the job every 10 minutes.

**Answer:** A

**Explanation:**
AWS Lambda supports container images as a packaging format for functions. You can use existing container development workflows to package and deploy Lambda functions as container images of up to 10 GB in size. You can also use familiar tools such as Docker CLI to build, test, and push your container images to Amazon Elastic Container Registry (Amazon ECR). You can then create an AWS Lambda function based on the container image of your job and configure Amazon EventBridge to invoke the function every 10 minutes using a cron expression. This solution will be cost-effective as you only pay for the compute time you consume when your function runs. References: https://docs.aws.amazon.com/lambda/latest/dg/images-create.html
https://docs.aws.amazon.com/eventbridge/latest/userguide/run-lambda-schedule.html

**NEW QUESTION 7**
- (Topic 4)
A social media company runs its application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. The application has more than a billion images stored in an Amazon S3 bucket and processes thousands of images each second. The company wants to resize the images dynamically and serve appropriate formats to clients.
Which solution will meet these requirements with the LEAST operational overhead?

A. Install an external image management library on an EC2 instanc
B. Use the image management library to process the images.
C. Create a CloudFront origin request polic
D. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.
E. Use a Lambda@Edge function with an external image management librar
F. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.
G. Create a CloudFront response headers polic
H. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.

**Answer:** C

**Explanation:**
Lambda@Edge is a service that allows you to run Lambda functions at CloudFront edge locations. It can be used to modify requests and responses that flow through CloudFront. CloudFront origin request policy is a policy that controls the values (URL query strings, HTTP headers, and cookies) that are included in requests that CloudFront sends to the origin. It can be used to collect additional information at the origin or to customize the origin response. CloudFront response headers policy is a policy that specifies the HTTP headers that CloudFront removes or adds in responses that it sends to viewers. It can be used to add security or custom headers to responses.
Based on these definitions, the solution that will meet the requirements with the least operational overhead is:
* C. Use a Lambda@Edge function with an external image management library. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.
This solution would allow the application to use a Lambda@Edge function to resize the images dynamically and serve appropriate formats to clients based on the User-Agent HTTP header in the request. The Lambda@Edge function would run at the edge locations,
reducing latency and load on the origin. The application code would only need to include an external image management library that can perform image manipulation tasks1.

**NEW QUESTION 8**
- (Topic 4)
An image hosting company uploads its large assets to Amazon S3 Standard buckets The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again. For the first 30 days after upload, the objects will be accessed frequently. The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.
Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

A. Move assets to S3 Intelligent-Tiering after 30 days.
B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.
C. Configure an S3 Lifecycle policy to clean up expired object delete markers.
D. Move assets to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days
E. Move assets to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

**Answer:** AB

**Explanation:**
S3 Intelligent-Tiering is a storage class that automatically moves data to the most cost-effective access tier based on access frequency, without performance impact, retrieval fees, or operational overhead1. It is ideal for data with unknown or changing access patterns, such as the company's assets. By moving assets to S3 Intelligent-Tiering after 30 days, the company can optimize its storage costs while maintaining high availability and resilience of stored assets.
S3 Lifecycle is a feature that enables you to manage your objects so that they are stored cost effectively throughout their lifecycle2. You can create lifecycle rules to define actions that Amazon S3 applies to a group of objects. One of the actions is to abort incomplete multipart uploads that can occur when an upload is interrupted. By configuring an S3 Lifecycle policy to clean up incomplete multipart uploads, the company can reduce its storage costs and avoid paying for parts that are not used.
Option C is incorrect because expired object delete markers are automatically deleted by Amazon S3 and do not incur any storage costs3. Therefore, configuring an S3 Lifecycle policy to clean up expired object delete markers will not have any effect on the company's storage costs.
Option D is incorrect because S3 Standard-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed1. It has a lower storage cost than S3 Standard, but it has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 Standard-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally.
Option E is incorrect because S3 One Zone-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed1. It has a lower storage cost than S3 Standard-IA, but it stores data in only one Availability Zone and has less resilience than other storage classes. It also has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 One Zone-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally or require high availability. Reference URL: 1: https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html 2:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html 3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/delete-or-empty- bucket.html#delete-bucket-considerations : https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html :
https://aws.amazon.com/certification/certified-solutions-architect-associate/

**NEW QUESTION 9**
- (Topic 4)
A company has users all around the world accessing its HTTP-based application deployed on Amazon EC2 instances in multiple AWS Regions. The company wants to improve the availability and performance of the application. The company also wants to protect the application against common web exploits that may affect availability, compromise security, or consume excessive resources. Static IP addresses are required.
What should a solutions architect recommend to accomplish this?

A. Put the EC2 instances behind Network Load Balancers (NLBs) in each Regio
B. Deploy AWS WAF on the NLB
C. Create an accelerator using AWS Global Accelerator and register the NLBs as endpoints.
D. Put the EC2 instances behind Application Load Balancers (ALBs) in each Regio

E. Deploy AWS WAF on the ALB
F. Create an accelerator using AWS Global Accelerator and register the ALBs as endpoints.
G. Put the EC2 instances behind Network Load Balancers (NLBs) in each Regio
H. Deploy AWS WAF on the NLB
I. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the NLBs.
J. Put the EC2 instances behind Application Load Balancers (ALBs) in each Regio
K. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the ALB
L. Deploy AWS WAF on the CloudFront distribution.

**Answer:** A

**Explanation:**
The company wants to improve the availability and performance of the application, as well as protect it against common web exploits. The company also needs static IP addresses for the application. To meet these requirements, a solutions architect should recommend the following solution:
? Put the EC2 instances behind Network Load Balancers (NLBs) in each Region.
NLBs are designed to handle millions of requests per second while maintaining high throughput at ultra-low latency. NLBs also support static IP addresses for each Availability Zone, which can be useful for whitelisting or firewalling purposes.
? Deploy AWS WAF on the NLBs. AWS WAF is a web application firewall that helps
protect web applications from common web exploits that could affect availability, security, or performance. AWS WAF lets you define customizable web security rules that control which traffic to allow or block to your web applications.
? Create an accelerator using AWS Global Accelerator and register the NLBs as
endpoints. AWS Global Accelerator is a service that improves the availability and performance of your applications with local or global users. It provides static IP addresses that act as a fixed entry point to your application endpoints in any AWS Region. It uses the AWS global network to optimize the path from your users to your applications, improving the performance of your TCP and UDP traffic.
This solution will provide high availability across Availability Zones and Regions, improve performance by routing traffic over the AWS global network, protect the application from common web attacks, and provide static IP addresses for the application.
References:
? Network Load Balancer
? AWS WAF
? AWS Global Accelerator

**NEW QUESTION 10**
- (Topic 4)
A company is creating an application The company stores data from tests of the application in multiple on-premises locations
The company needs to connect the on-premises locations to VPCs in an AWS Region in the AWS Cloud The number of accounts and VPCs will increase during the next year The network architecture must simplify the administration of new connections and must provide the ability to scale.
Which solution will meet these requirements with the LEAST administrative overhead'?

A. Create a peering connection between the VPCs Create a VPN connection between the VPCs and the on-premises locations.
B. Launch an Amazon EC2 instance On the instance, include VPN software that uses a VPN connection to connect all VPCs and on-premises locations.
C. Create a transit gateway Create VPC attachments for the VPC connections Create VPN attachments for the on-premises connections.
D. Create an AWS Direct Connect connection between the on-premises locations and acentral VP
E. Connect the central VPC to other VPCs by using peering connections.

**Answer:** C

**Explanation:**
A transit gateway is a network transit hub that enables you to connect your VPCs and on-premises networks in a centralized and scalable way. You can create VPC attachments to connect your VPCs to the transit gateway, and VPN attachments to connect your on-premises networks to the transit gateway over the internet. The transit gateway acts as a router between the attached networks, and simplifies the administration of new connections by reducing the number of peering or VPN connections required. You can also use transit gateway route tables to control the routing of traffic between the attached networks. By creating a transit gateway and using VPC and VPN attachments, you can meet the requirements of the company with the least administrative overhead.
References:
? AWS Transit Gateway
? Transit gateway attachments
? Transit gateway route tables

**NEW QUESTION 10**
- (Topic 4)
A company maintains an Amazon RDS database that maps users to cost centers. The company has accounts in an organization in AWS Organizations. The company needs a solution that will tag all resources that are created in a specific AWS account in the organization. The solution must tag each resource with the cost center ID of the user who created the resource.
Which solution will meet these requirements?

A. Move the specific AWS account to a new organizational unit (OU) in Organizations from the management accoun
B. Create a service control policy (SCP) that requires all existing resources to have the correct cost center tag before the resources are create
C. Apply the SCP to the new OU.
D. Create an AWS Lambda function to tag the resources after the Lambda function looks up the appropriate cost center from the RDS databas
E. Configure an Amazon EventBridge rule that reacts to AWS CloudTrail events to invoke the Lambda function.
F. Create an AWS CloudFormation stack to deploy an AWS Lambda functio
G. Configure the Lambda function to look up the appropriate cost center from the RDS database and to tag resource
H. Create an Amazon EventBridge scheduled rule to invoke the CloudFormation stack.
I. Create an AWS Lambda function to tag the resources with a default valu
J. Configure an Amazon EventBridge rule that reacts to AWS CloudTrail events to invoke the Lambda function when a resource is missing the cost center tag.

**Answer:** B

**Explanation:**
AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be used to tag resources with the cost center ID of the user who created the resource, by querying the RDS database that maps users to cost centers. Amazon EventBridge is a serverless event bus service that enables event-driven architectures. EventBridge can be configured to react to AWS CloudTrail events, which are recorded API calls made by or on

behalf of the AWS account. EventBridge can invoke the Lambda function when a resource is created in the specific AWS account, passing the user identity and resource information as parameters. This solution will meet the requirements, as it enables automatic tagging of resources based on the user and cost center mapping.
References:
? 1 provides an overview of AWS Lambda and its benefits.
? 2 provides an overview of Amazon EventBridge and its benefits.
? 3 explains the concept and benefits of AWS CloudTrail events.

**NEW QUESTION 15**
- (Topic 4)
A company runs an application on AWS. The application receives inconsistent amounts of usage. The application uses AWS Direct Connect to connect to an on-premises MySQL- compatible database. The on-premises database consistently uses a minimum of 2 GiB of memory.
The company wants to migrate the on-premises database to a managed AWS service. The company wants to use auto scaling capabilities to manage unexpected workload increases.
Which solution will meet these requirements with the LEAST administrative overhead?

A. Provision an Amazon DynamoDB database with default read and write capacity settings.
B. Provision an Amazon Aurora database with a minimum capacity of 1 Aurora capacityunit (ACU).
C. Provision an Amazon Aurora Serverless v2 database with a minimum capacity of 1 Aurora capacity unit (ACU).
D. Provision an Amazon RDS for MySQL database with 2 GiB of memory.

**Answer:** C

**Explanation:**
it allows the company to migrate the on-premises database to a managed AWS service that supports auto scaling capabilities and has the least administrative overhead. Amazon Aurora Serverless v2 is a configuration of Amazon Aurora that automatically scales compute capacity based on workload demand. It can scale from hundreds to hundreds of thousands of transactions in a fraction of a second. Amazon Aurora Serverless v2 also supports MySQL-compatible databases and AWS Direct Connect connectivity. References:
? Amazon Aurora Serverless v2
? Connecting to an Amazon Aurora DB Cluster

**NEW QUESTION 20**
- (Topic 4)
A company offers a food delivery service that is growing rapidly. Because of the growth, the company's order processing system is experiencing scaling problems during peak traffic hours. The current architecture includes the following:
• A group of Amazon EC2 instances that run in an Amazon EC2 Auto Scaling group to collect orders from the application
• Another group of EC2 instances that run in an Amazon EC2 Auto Scaling group to fulfill orders
The order collection process occurs quickly, but the order fulfillment process can take longer. Data must not be lost because of a scaling event.
A solutions architect must ensure that the order collection process and the order fulfillment process can both scale properly during peak traffic hours. The solution must optimize
utilization of the company's AWS resources. Which solution meets these requirements?

A. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
B. Configure each Auto Scaling group's minimum capacity according to peak workload values.
C. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
D. Configure a CloudWatch alarm to invoke an Amazon Simple Notification Service (Amazon SNS) topic that creates additional Auto Scaling groups on demand.
E. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillmen
F. Configure the EC2 instances to poll their respective queu
G. Scale the Auto Scaling groups based on notifications that the queues send.
H. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillmen
I. Configure the EC2 instances to poll their respective queu
J. Create a metric based on a backlog per instance calculatio
K. Scale the Auto Scaling groups based on this metric.

**Answer:** D

**Explanation:**
The number of instances in your Auto Scaling group can be driven by how long it takes to process a message and the acceptable amount of latency (queue delay). The solution is to use a backlog per instance metric with the target value being the acceptable backlog per instance to maintain.

**NEW QUESTION 21**
- (Topic 4)
A company wants to move from many standalone AWS accounts to a consolidated, multi- account architecture The company plans to create many new AWS accounts for different business units. The company needs to authenticate access to these AWS accounts by using a centralized corporate directory service.
Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

A. Create a new organization in AWS Organizations with all features turned o
B. Create the new AWS accounts in the organization.
C. Set up an Amazon Cognito identity poo
D. Configure AWS 1AM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication.
E. Configure a service control policy (SCP) to manage the AWS account
F. Add AWS 1AM Identity Center (AWS Single Sign-On) to AWS Directory Service.
G. Create a new organization in AWS Organization
H. Configure the organization's authentication mechanism to use AWS Directory Service directly.
I. Set up AWS 1AM Identity Center (AWS Single Sign-On) in the organizatio
J. Configure 1AM Identity Center, and integrate it with the company's corporate directory service.

**Answer:** AE

**Explanation:**

AWS Organizations is a service that helps users centrally manage and govern multiple AWS accounts. It allows users to create organizational units (OUs) to group accounts based on business needs or other criteria. It also allows users to define and attach service control policies (SCPs) to OUs or accounts to restrict the actions that can be performed by the accounts1. By creating a new organization in AWS Organizations with all features turned on, the solution can consolidate and manage the new AWS accounts for different business units.
AWS IAM Identity Center (formerly known as AWS Single Sign-On) is a service that provides single sign-on access for all of your AWS accounts and cloud applications. It connects with Microsoft Active Directory through AWS Directory Service to allow users in that directory to sign in to a personalized AWS access portal using their existing Active Directory user names and passwords. From the AWS access portal, users have access to all the AWS accounts and cloud applications that they have permissions for2. By setting up IAM Identity Center in the organization and integrating it with the company's corporate directory service, the solution can authenticate access to these AWS accounts using a centralized corporate directory service.
* B. Set up an Amazon Cognito identity pool. Configure AWS 1AM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication. This solution will not meet the requirement of authenticating access to these AWS accounts by using a centralized corporate directory service, as Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications, not for corporate directory services3.
* C. Configure a service control policy (SCP) to manage the AWS accounts. Add AWS 1AM Identi-ty Center (AWS Single Sign-On) to AWS Directory Service. This solution will not work, as SCPs are used to restrict the actions that can be performed by the accounts in an organization, not to manage the accounts themselves1. Also, IAM Identity Center cannot be added to AWS Directory Service, as it is a separate service that connects with Microsoft Active Directory through AWS Directory Service2.
* D. Create a new organization in AWS Organizations. Configure the organization's authentication mechanism to use AWS Directory Service directly. This solution will not work, as AWS Organizations does not have an authentication mechanism that can use AWS Directory Service directly. AWS Organizations relies on IAM Identity Center to provide single sign-on access for the accounts in an organization.
Reference URL: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_integrate_services.html

**NEW QUESTION 25**
- (Topic 4)
An online video game company must maintain ultra-low latency for its game servers. The game servers run on Amazon EC2 instances. The company needs a solution that can
handle millions of UDP internet traffic requests each second.
Which solution will meet these requirements MOST cost-effectively?

A. Configure an Application Load Balancer with the required protocol and ports for the internet traffi
B. Specify the EC2 instances as the targets.
C. Configure a Gateway Load Balancer for the internet traffi
D. Specify the EC2 instances as the targets.
E. Configure a Network Load Balancer with the required protocol and ports for the internet traffi
F. Specify the EC2 instances as the targets.
G. Launch an identical set of game servers on EC2 instances in separate AWS Region
H. Route internet traffic to both sets of EC2 instances.

**Answer:** C

**Explanation:**
The most cost-effective solution for the online video game company is to configure a Network Load Balancer with the required protocol and ports for the internet traffic and specify the EC2 instances as the targets. This solution will enable the company to handle millions of UDP requests per second with ultra-low latency and high performance. A Network Load Balancer is a type of Elastic Load Balancing that operates at the connection level (Layer 4) and routes traffic to targets (EC2 instances, microservices, or containers) within Amazon VPC based on IP protocol data. A Network Load Balancer is ideal for load balancing of both TCP and UDP traffic, as it is capable of handling millions of requests per second while maintaining high throughput at ultra-low latency. A Network Load Balancer also preserves the source IP address of the clients to the back-end applications, which can be useful for logging or security purposes1.

**NEW QUESTION 29**
- (Topic 4)
A company uses Amazon EC2 instances to host its internal systems. As part of a deployment operation, an administrator tries to use the AWS CLI to terminate an EC2 instance. However, the administrator receives a 403 (Access Denied) error message.
The administrator is using an IAM role that has the following IAM policy attached:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": ["ec2:TerminateInstances"],
            "Resource": ["*"]
        },
        {
            "Effect": "Deny",
            "Action": ["ec2:TerminateInstances"],
            "Condition": {
                "NotIpAddress": {
                    "aws:SourceIp": [
                        "192.0.2.0/24",
                        "203.0.113.0/24"
                    ]
                },
                "Resource": ["*"]
            }
    ]
}
```

What is the cause of the unsuccessful request?

A. The EC2 instance has a resource-based policy with a Deny statement.
B. The principal has not been specified in the policy statement
C. The "Action" field does not grant the actions that are required to terminate the EC2 instance.
D. The request to terminate the EC2 instance does not originate from the CIDR blocks 192.0.2.0/24 or 203.0 113.0/24

**Answer:** D

**NEW QUESTION 30**
- (Topic 4)
A company wants to migrate its three-tier application from on premises to AWS. The web tier and the application tier are running on third-party virtual machines (VMs). The database tier is running on MySQL.
The company needs to migrate the application by making the fewest possible changes to the architecture. The company also needs a database solution that can restore data to a specific point in time.
Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate the web tier and the application tier to Amazon EC2 instances in private subnet
B. Migrate the database tier to Amazon RDS for MySQL in private subnets.
C. Migrate the web tier to Amazon EC2 instances in public subnet
D. Migrate the application tier to EC2 instances in private subnet
E. Migrate the database tier to Amazon Aurora MySQL in private subnets.
F. Migrate the web tier to Amazon EC2 instances in public subnet
G. Migrate the application tier to EC2 instances in private subnet
H. Migrate the database tier to Amazon RDS for MySQL in private subnets.
I. Migrate the web tier and the application tier to Amazon EC2 instances in public subnet
J. Migrate the database tier to Amazon Aurora MySQL in public subnets.

**Answer:** C

**Explanation:**
The solution that meets the requirements with the least operational overhead is to migrate the web tier to Amazon EC2 instances in public subnets, migrate the application tier to EC2 instances in private subnets, and migrate the database tier to Amazon RDS for MySQL in private subnets. This solution allows the company to migrate its three-tier application to AWS by making minimal changes to the architecture, as it preserves the same web, application, and database tiers and uses the same MySQL database engine. The solution also provides a database solution that can restore data to a specific point in time, as Amazon RDS for MySQL supports automated backups and point- in-time recovery. This solution also reduces the operational overhead by using managed services such as Amazon EC2 and Amazon RDS, which handle tasks such as provisioning, patching, scaling, and monitoring.
The other solutions do not meet the requirements as well as the first one because they either involve more changes to the architecture, do not provide point-in-time recovery, or do not follow best practices for security and availability. Migrating the database tier to Amazon Aurora MySQL would require changing the database engine and potentially modifying the application code to ensure compatibility. Migrating the web tier and the application tier to public subnets would expose them to more security risks and reduce their availability in case of a subnet failure. Migrating the database tier to public subnets would also compromise its security and performance. References:
? Migrate Your Application Database to Amazon RDS
? Amazon RDS for MySQL
? Amazon Aurora MySQL
? Amazon VPC

**NEW QUESTION 33**
- (Topic 4)
A solutions architect is implementing a document review application using an Amazon S3 bucket for storage. The solution must prevent accidental deletion of the documents and ensure that all versions of the documents are available. Users must be able to download, modify, and upload documents.
Which combination of actions should be taken to meet these requirements? (Choose two.)

A. Enable a read-only bucket ACL.
B. Enable versioning on the bucket.
C. Attach an IAM policy to the bucket.
D. Enable MFA Delete on the bucket.
E. Encrypt the bucket using AWS KMS.

**Answer:** BD

**Explanation:**
Versioning is a feature of Amazon S3 that allows users to keep multiple versions of the same object in a bucket. It can help prevent accidental deletion of the documents and ensure that all versions of the documents are available1. MFA Delete is a feature of Amazon S3 that adds an extra layer of security by requiring two forms of authentication to delete a version or change the versioning state of a bucket. It can help prevent unauthorized or accidental deletion of the documents2. By enabling both versioning and MFA Delete on the bucket, the solution can meet the requirements.
* A. Enable a read-only bucket ACL. This solution will not meet the requirement of allowing
users to download, modify, and upload documents, as a read-only bucket ACL will prevent write access to the bucket3.
* C. Attach an IAM policy to the bucket. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as an IAM policy is used to grant or deny permissions to users or roles, not to enable versioning or MFA Delete4.
* E. Encrypt the bucket using AWS KMS. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as encrypting the bucket using AWS KMS is a method of protecting data at rest, not enabling versioning or MFA Delete.
Reference URL: https://docs.aws.amazon.com/AmazonS3/latest/userguide/Versioning.html

**NEW QUESTION 38**
- (Topic 4)
A recent analysis of a company's IT expenses highlights the need to reduce backup costs. The company's chief information officer wants to simplify the on-premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. The company must preserve the existing investment in the on-premises backup applications and workflows.
What should a solutions architect recommend?

A. Set up AWS Storage Gateway to connect with the backup applications using the NFS interface.
B. Set up an Amazon EFS file system that connects with the backup applications using the NFS interface.
C. Set up an Amazon EFS file system that connects with the backup applications using the iSCSI interface.
D. Set up AWS Storage Gateway to connect with the backup applications using the iSCSI- virtual tape library (VTL) interface.

**Answer:** D

**Explanation:**
it allows the company to simplify the on-premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. By setting up AWS Storage Gateway to connect with the backup applications using the iSCSI-virtual tape library (VTL) interface, the company can store backup data on virtual tapes in S3 or Glacier. This preserves the existing investment in the on-premises backup applications and workflows while leveraging AWS storage services.
References:
? AWS Storage Gateway
? Tape Gateway

**NEW QUESTION 40**
- (Topic 4)
A company has a mobile chat application with a data store based in Amazon uynamoUb. users would like new messages to be read with as little latency as possible A solutions architect needs to design an optimal solution that requires minimal application changes.
Which method should the solutions architect select?

A. Configure Amazon DynamoDB Accelerator (DAX) for the new messages tabl
B. Update the code to use the DAXendpoint.
C. Add DynamoDB read repticas to handle the increased read loa
D. Update the application to point to the read endpoint for the read replicas.
E. Double the number of read capacity units for the new messages table in DynamoD
F. Continue to use the existing DynamoDB endpoint.
G. Add an Amazon ElastiCache for Redis cache to the application stac
H. Update the application to point to the Redis cache endpoint instead of DynamoDB.

**Answer:** A

**Explanation:**
https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-high- latency/
Amazon DynamoDB Accelerator (DAX) is a fully managed in-memory cache for DynamoDB that improves the performance of DynamoDB tables by up to 10 times and
provides microsecond level of response time at any scale. It is compatible with DynamoDB API operations and requires minimal code changes to use1. By configuring DAX for the
new messages table, the solution can reduce the latency for reading new messages with minimal application changes.
* B. Add DynamoDB read repticas to handle the increased read load. Update the application to point to the read endpoint for the read replicas. This solution will not work, as DynamoDB does not support read replicas as a feature. Read replicas are available for Amazon RDS, not for DynamoDB2.
* C. Double the number of read capacity units for the new messages table in DynamoDB. Continue to use the existing DynamoDB endpoint. This solution will not meet the requirement of reading new messages with as little latency as possible, as increasing the
read capacity units will only increase the throughput of DynamoDB, not the performance or latency3.
* D. Add an Amazon ElastiCache for Redis cache to the application stack. Update the application to point to the Redis cache endpoint instead of DynamoDB. This solution will not meet the requirement of minimal application changes, as adding ElastiCache for Redis will require significant code changes to implement caching logic, such as querying cache first, updating cache after writing to DynamoDB, and invalidating cache when needed. Reference URL:

https://aws.amazon.com/dynamodb/dax/

**NEW QUESTION 45**
- (Topic 4)
A company used an Amazon RDS for MySQL DB instance during application testing. Before terminating the DB instance at the end of the test cycle, a solutions architect created two backups. The solutions architect created the first backup by using the mysqldump utility to create a database dump. The solutions architect created the second backup by enabling the final DB snapshot option on RDS termination.
The company is now planning for a new test cycle and wants to create a new DB instance from the most recent backup. The company has chosen a MySQL-compatible edition of Amazon Aurora to host the DB instance.
Which solutions will create the new DB instance? (Select TWO.)

A. Import the RDS snapshot directly into Aurora.
B. Upload the RDS snapshot to Amazon S3. Then import the RDS snapshot into Aurora.
C. Upload the database dump to Amazon S3. Then import the database dump into Aurora.
D. Use AWS Database Migration Service (AWS DMS) to import the RDS snapshot into Aurora.
E. Upload the database dump to Amazon S3. Then use AWS Database Migration Service (AWS DMS) to import the database dump into Aurora.

**Answer:** AC

**Explanation:**
These answers are correct because they meet the requirements of creating a new DB instance from the most recent backup and using a MySQL-compatible edition of Amazon Aurora to host the DB instance. You can import the RDS snapshot directly into Aurora if the MySQL DB instance and the Aurora DB cluster are running the same version of MySQL. For example, you can restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.6, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is simple and requires the fewest number of steps. You can upload the database dump to Amazon S3 and then import the database dump into Aurora if the MySQL DB instance and the Aurora DB cluster are running different versions of MySQL. For example, you can import a MySQL version 5.6 database dump into Aurora MySQL version 5.7, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is more flexible and allows you to migrate across different versions of MySQL.
References:
? https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL
.Migrating.RDSMySQL.Import.html
? https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL
.Migrating.RDSMySQL.Dump.html

**NEW QUESTION 50**
- (Topic 4)
A company stores multiple Amazon Machine Images (AMIs) in an AWS account to launch its Amazon EC2 instances. The AMIs contain critical data and configurations that are necessary for the company's operations. The company wants to implement a solution that will recover accidentally deleted AMIs quickly and efficiently.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create Amazon Elastic Block Store (Amazon EBS) snapshots of the AMI
B. Store the snapshots in a separate AWS account.
C. Copy all AMIs to another AWS account periodically.
D. Create a retention rule in Recycle Bin.
E. Upload the AMIs to an Amazon S3 bucket that has Cross-Region Replication.

**Answer:** C

**Explanation:**
Recycle Bin is a data recovery feature that enables you to restore accidentally deleted Amazon EBS snapshots and EBS-backed AMIs. When using Recycle Bin, if your resources are deleted, they are retained in the Recycle Bin for a time period that you specify before being permanently deleted. You can restore a resource from the Recycle Bin at any time before its retention period expires. This solution has the least operational overhead, as you do not need to create, copy, or upload any additional resources. You can also manage tags and permissions for AMIs in the Recycle Bin. AMIs in the Recycle Bin do not incur any additional charges.
References:
? Recover AMIs from the Recycle Bin
? Recover an accidentally deleted Linux AMI

**NEW QUESTION 51**
- (Topic 4)
A company runs an infrastructure monitoring service. The company is building a new feature that will enable the service to monitor data in customer AWS accounts. The new feature will call AWS APIs in customer accounts to describe Amazon EC2 instances and read Amazon CloudWatch metrics.
What should the company do to obtain access to customer accounts in the MOST secure way?

A. Ensure that the customers create an 1AM role in their account with read-only EC2 and CloudWatch permissions and a trust policy to the company's account.
B. Create a serverless API that implements a token vending machine to provide temporary AWS credentials for a role with read-only EC2 and CloudWatch permissions.
C. Ensure that the customers create an 1AM user in their account with read-only EC2 and CloudWatch permission
D. Encrypt and store customer access and secret keys in a secrets management system.
E. Ensure that the customers create an Amazon Cognito user in their account to use an 1AM role with read-only EC2 and CloudWatch permission
F. Encrypt and store the Amazon Cognito user and password in a secrets management system.

**Answer:** A

**Explanation:**
By having customers create an IAM role with the necessary permissions in their own accounts, the company can use AWS Identity and Access Management (IAM) to establish cross-account access. The trust policy allows the company's AWS account to assume the customer's IAM role temporarily, granting access to the specified resources (EC2 instances and CloudWatch metrics) within the customer's account. This approach follows the principle of least privilege, as the company only requests the necessary permissions and does not require long-term access keys or user credentials from the customers.

**NEW QUESTION 55**
- (Topic 4)
A company has deployed a multiplayer game for mobile devices. The game requires live
location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations.
The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly.
What should a solutions architect do to improve the performance of the data tier?

A. Take a snapshot of the existing DB instanc
B. Restore the snapshot with Multi-AZ enabled.
C. Migrate from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards.
D. Deploy Amazon DynamoDB Accelerator (DAX) in front of the existing DB instanc
E. Modify the game to use DAX.
F. Deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instanc
G. Modify the game to use Redis.

**Answer:** D

**Explanation:**
The solution that will improve the performance of the data tier is to deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance and modify the game to use Redis. This solution will enable the game to store and retrieve the location data of the players in a fast and scalable way, as Redis is an in-memory data store that supports geospatial data types and commands. By using ElastiCache for Redis, the game can reduce the load on the RDS for PostgreSQL DB instance, which is not optimized for high-frequency updates and queries of location data. ElastiCache for Redis also supports replication, sharding, and auto scaling to handle the increasing user base of the game. The other solutions are not as effective as the first one because they either do not improve the performance, do not support geospatial data, or do not leverage caching. Taking a snapshot of the existing DB instance and restoring it with Multi-AZ enabled will not improve the performance of the data tier, as it only provides high availability and durability, but not scalability or low latency. Migrating from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards will not improve the performance of the data tier, as OpenSearch Service is mainly designed for full-text search and analytics, not for real-time location tracking. OpenSearch Service also does not support geospatial data types and commands natively, unlike Redis. Deploying Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance and modifying the game to use DAX will not improve the performance of the data tier, as DAX is only compatible with DynamoDB, not with RDS for PostgreSQL. DAX also does not support geospatial data types and commands.
References:
? Amazon ElastiCache for Redis
? Geospatial Data Support - Amazon ElastiCache for Redis
? Amazon RDS for PostgreSQL
? Amazon OpenSearch Service
? Amazon DynamoDB Accelerator (DAX)

**NEW QUESTION 58**
- (Topic 4)
A company runs demonstration environments for its customers on Amazon EC2 instances. Each environment is isolated in its own VPC. The company’s operations team needs to be notified when RDP or SSH access to an environment has been established.

A. Configure Amazon CloudWatch Application Insights to create AWS Systems Manager OpsItems when RDP or SSH access is detected.
B. Configure the EC2 instances with an IAM instance profile that has an IAM role with the AmazonSSMManagedInstanceCore policy attached.
C. Publish VPC flow logs to Amazon CloudWatch Log
D. Create required metric filter
E. Create an Amazon CloudWatch metric alarm with a notification action for when the alarm is in the ALARM state.
F. Configure an Amazon EventBridge rule to listen for events of type EC2 Instance State- change Notificatio
G. Configure an Amazon Simple Notification Service (Amazon SNS) topic as a targe
H. Subscribe the operations team to the topic.

**Answer:** C

**Explanation:**
https://aws.amazon.com/blogs/security/how-to-monitor-and-visualize-failed- ssh-access-attempts-to-amazon-ec2-linux-instances/

**NEW QUESTION 61**
- (Topic 4)
A company is concerned that two NAT instances in use will no longer be able to support the traffic needed for the company’s application. A solutions architect wants to implement a solution that is highly available, fault tolerant, and automatically scalable.
What should the solutions architect recommend?

A. Remove the two NAT instances and replace them with two NAT gateways in the same Availability Zone.
B. Use Auto Scaling groups with Network Load Balancers for the NAT instances in differentAvailability Zones.
C. Remove the two NAT instances and replace them with two NAT gateways in different Availability Zones.
D. Replace the two NAT instances with Spot Instances in different Availability Zones and deploy a Network Load Balancer.

**Answer:** C

**Explanation:**
If you have resources in multiple Availability Zones and they share one NAT gateway, and if the NAT gateway’s Availability Zone is down, resources in the other Availability Zones lose internet access. To create an Availability Zone-independent architecture, create a NAT gateway in each Availability Zone and configure your routing to ensure that resources use the NAT gateway in the same Availability Zone. https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html#nat-gateway- basics

**NEW QUESTION 66**
- (Topic 4)
A retail company uses a regional Amazon API Gateway API for its public REST APIs. The API Gateway endpoint is a custom domain name that points to an Amazon Route 53 alias record. A solutions architect needs to create a solution that has minimal effects on customers and minimal data loss to release the new version of APIs.

Which solution will meet these requirements?

A. Create a canary release deployment stage for API Gatewa
B. Deploy the latest API versio
C. Point an appropriate percentage of traffic to the canary stag
D. After API verification, promote the canary stage to the production stage.
E. Create a new API Gateway endpoint with a new version of the API in OpenAPI YAMLfile forma
F. Use the import-to-update operation in merge mode into the API in API Gatewa
G. Deploy the new version of the API to the production stage.
H. Create a new API Gateway endpoint with a new version of the API in OpenAPI JSON file forma
I. Use the import-to-update operation in overwrite mode into the API in API Gatewa
J. Deploy the new version of the API to the production stage.
K. Create a new API Gateway endpoint with new versions of the API definition
L. Create a custom domain name for the new API Gateway AP
M. Point the Route 53 alias record to the new API Gateway API custom domain name.

**Answer:** A

**Explanation:**
This answer is correct because it meets the requirements of releasing the new version of APIs with minimal effects on customers and minimal data loss. A canary release deployment is a software development strategy in which a new version of an API is deployed for testing purposes, and the base version remains deployed as a production release for normal operations on the same stage. In a canary release deployment, total API traffic is separated at random into a production release and a canary release with a pre- configured ratio. Typically, the canary release receives a small percentage of API traffic and the production release takes up the rest. The updated API features are only visible to API traffic through the canary. You can adjust the canary traffic percentage to optimize test coverage or performance. By keeping canary traffic small and the selection random, most users are not adversely affected at any time by potential bugs in the new version, and no single user is adversely affected all the time. After the test metrics pass your requirements, you can promote the canary release to the production release and disable the canary from the deployment. This makes the new features available in the production stage. References:
? https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html

**NEW QUESTION 71**
- (Topic 4)
A company is conducting an internal audit. The company wants to ensure that the data in an Amazon S3 bucket that is associated with the company's AWS Lake Formation data lake does not contain sensitive customer or employee data. The company wants to discover personally identifiable information (Pll) or financial information, including passport numbers and credit card numbers.
Which solution will meet these requirements?

A. Configure AWS Audit Manager on the accoun
B. Select the Payment Card Industry Data Security Standards (PCI DSS) for auditing.
C. Configure Amazon S3 Inventory on the S3 bucke
D. Configure Amazon Athena to query the inventory.
E. Configure Amazon Macie to run a data discovery job that uses managed identifiers for the required data types.
F. Use Amazon S3 Select to run a report across the S3 bucket.

**Answer:** C

**Explanation:**
Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS. Macie can run data discovery jobs that use managed identifiers for various types of PII or financial information, such as passport numbers and credit card numbers. Macie can also generate findings that alert you to potential issues or risks with your data. References:
https://docs.aws.amazon.com/macie/latest/userguide/macie-identifiers.html

**NEW QUESTION 73**
- (Topic 4)
A company wants to use high-performance computing and artificial intelligence to improve its fraud prevention and detection technology. The company requires distributed processing to complete a single workload as quickly as possible.
Which solution will meet these requirements?

A. Use Amazon Elastic Kubernetes Service (Amazon EKS) and multiple containers.
B. Use AWS ParallelCluster and the Message Passing Interface (MPI) libraries.
C. Use an Application Load Balancer and Amazon EC2 instances.
D. Use AWS Lambda functions.

**Answer:** B

**Explanation:**
AWS ParallelCluster is a service that allows you to create and manage high- performance computing (HPC) clusters on AWS. It supports multiple schedulers, including AWS Batch, which can run distributed workloads across multiple EC2 instances1.
MPI is a standard for message passing between processes in parallel computing. It provides functions for sending and receiving data, synchronizing processes, and managing communication groups2.
By using AWS ParallelCluster and MPI libraries, you can take advantage of the following benefits:
? You can easily create and configure HPC clusters that meet your specific requirements, such as instance type, number of nodes, network configuration, and storage options1.
? You can leverage the scalability and elasticity of AWS to run large-scale parallel
workloads without worrying about provisioning or managing servers1.
? You can use MPI libraries to optimize the performance and efficiency of your parallel applications by enabling inter-process communication and data exchange2.
? You can choose from a variety of MPI implementations that are compatible with AWS ParallelCluster, such as Open MPI, Intel MPI, and MPICH3.

**NEW QUESTION 78**
- (Topic 4)
A company operates an ecommerce website on Amazon EC2 instances behind an Application Load Balancer (ALB) in an Auto Scaling group. The site is

experiencing performance issues related to a high request rate from illegitimate external systems with changing IP addresses. The security team is worried about potential DDoS attacks against the website. The company must block the illegitimate incoming requests in a way that has a minimal impact on legitimate users. What should a solutions architect recommend?

A. Deploy Amazon Inspector and associate it with the ALB.
B. Deploy AWS WAF, associate it with the ALB, and configure a rate-limiting rule.
C. Deploy rules to the network ACLs associated with the ALB to block the incoming traffic.
D. Deploy Amazon GuardDuty and enable rate-limiting protection when configuring GuardDuty.

**Answer:** B

**Explanation:**
This answer is correct because it meets the requirements of blocking the illegitimate incoming requests in a way that has a minimal impact on legitimate users. AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define. You can associate AWS WAF with an ALB to protect the web application from malicious requests. You can configure a rate-limiting rule in AWS WAF to track the rate of requests for each originating IP address and block requests from an IP address that exceeds a certain limit within a five-minute period. This way, you can mitigate potential DDoS attacks and improve the performance of your website.
References:
? https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html
? https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type- rate-based.html

**NEW QUESTION 80**
- (Topic 4)
A company sends AWS CloudTrail logs from multiple AWS accounts to an Amazon S3 bucket in a centralized account. The company must keep the CloudTrail logs. The company must also be able to query the CloudTrail logs at any time
Which solution will meet these requirements?

A. Use the CloudTraiI event history in the centralized account to create an Amazon Athena tabl
B. Query the CloudTrail logs from Athena.
C. Configure an Amazon Neptune instance to manage the CloudTrail log
D. Query the CloudTraiI logs from Neptune.
E. Configure CloudTrail to send the logs to an Amazon DynamoDB tabl
F. Create a dashboard in Amazon QulCkSight to query the logs in the table.
G. use Amazon Athena to create an Athena noteboo
H. Configure CloudTrail to send the logs to the noteboo
I. Run queries from Athena.

**Answer:** A

**Explanation:**
it allows the company to keep the CloudTrail logs and query them at any time. By using the CloudTrail event history in the centralized account, the company can view, filter, and download recent API activity across multiple AWS accounts. By creating an Amazon Athena table from the CloudTrail event history, the company can use a serverless interactive query service that makes it easy to analyze data in S3 using standard SQL. By querying the CloudTrail logs from Athena, the company can gain insights into user activity and resource changes. References:
? Viewing Events with CloudTrail Event History
? Querying AWS CloudTrail Logs
? Amazon Athena

**NEW QUESTION 85**
- (Topic 4)
A company is deploying an application that processes streaming data in near-real time The company plans to use Amazon EC2 instances for the workload The network architecture must be configurable to provide the lowest possible latency between nodes
Which combination of network solutions will meet these requirements? (Select TWO)

A. Enable and configure enhanced networking on each EC2 instance
B. Group the EC2 instances in separate accounts
C. Run the EC2 instances in a cluster placement group
D. Attach multiple elastic network interfaces to each EC2 instance
E. Use Amazon Elastic Block Store (Amazon EBS) optimized instance types.

**Answer:** AC

**Explanation:**
These options are the most suitable ways to configure the network architecture to provide the lowest possible latency between nodes. Option A enables and configures enhanced networking on each EC2 instance, which is a feature that improves the network performance of the instance by providing higher bandwidth, lower latency, and lower jitter. Enhanced networking uses single root I/O virtualization (SR-IOV) or Elastic Fabric Adapter (EFA) to provide direct access to the network hardware. You can enable and configure enhanced networking by choosing a supported instance type and a compatible operating system, and installing the required drivers. Option C runs the EC2 instances in a cluster placement group, which is a logical grouping of instances within a single Availability Zone that are placed close together on the same underlying hardware. Cluster placement groups provide the lowest network latency and the highest network throughput among the placement group options. You can run the EC2 instances in a cluster placement group by creating a placement group and launching the instances into it. Option B is not suitable because grouping the EC2 instances in separate accounts does not provide the lowest possible latency between nodes. Separate accounts are used to isolate and organize resources for different purposes, such as security, billing, or compliance. However, they do not affect the network performance or proximity of the instances. Moreover, grouping the EC2 instances in separate accounts would incur additional costs and complexity, and it would require setting up cross-account networking and permissions.
Option D is not suitable because attaching multiple elastic network interfaces to each EC2 instance does not provide the lowest possible latency between nodes. Elastic network interfaces are virtual network interfaces that can be attached to EC2 instances to provide additional network capabilities, such as multiple IP addresses, multiple subnets, or enhanced security. However, they do not affect the network performance or proximity of the instances. Moreover, attaching multiple elastic network interfaces to each EC2 instance would consume additional resources and limit the instance type choices.
Option E is not suitable because using Amazon EBS optimized instance types does not provide the lowest possible latency between nodes. Amazon EBS

optimized instance types are instances that provide dedicated bandwidth for Amazon EBS volumes, which are block storage volumes that can be attached to EC2 instances. EBS optimized instance types improve the performance and consistency of the EBS volumes, but they do not affect the network performance or proximity of the instances. Moreover, using EBS optimized instance types would incur additional costs and may not be necessary for the streaming data workload.
References:
? Enhanced networking on Linux
? Placement groups
? Elastic network interfaces
? Amazon EBS-optimized instances

**NEW QUESTION 87**
- (Topic 4)
A company has created a multi-tier application for its ecommerce website. The website uses an Application Load Balancer that resides in the public subnets, a web tier in the public subnets, and a MySQL cluster hosted on Amazon EC2 instances in the private subnets. The MySQL database needs to retrieve product catalog and pricing information that is hosted on the internet by a third-party provider. A solutions architect must devise a strategy that maximizes security without increasing operational overhead. What should the solutions architect do to meet these requirements?

A. Deploy a NAT instance in the VP
B. Route all the internet-based traffic through the NAT instance.
C. Deploy a NAT gateway in the public subnet
D. Modify the private subnet route table to direct all internet-bound traffic to the NAT gateway.
E. Configure an internet gateway and attach it to the VP
F. Modify the private subnet route table to direct internet-bound traffic to the internet gateway.
G. Configure a virtual private gateway and attach it to the VP
H. Modify the private subnet route table to direct internet-bound traffic to the virtual private gateway.

**Answer:** B

**Explanation:**
To allow the MySQL database in the private subnets to access the internet without exposing it to the public, a NAT gateway is a suitable solution. A NAT gateway enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances. A NAT gateway resides in the public subnets and can handle high throughput of traffic with low latency. A NAT gateway is also a managed service that does not require any operational overhead. References:
? NAT Gateways
? NAT Gateway Pricing

**NEW QUESTION 88**
- (Topic 4)
A company wants to manage Amazon Machine Images (AMIs). The company currently copies AMIs to the same AWS Region where the AMIs were created. The company needs to design an application that captures AWS API calls and sends alerts whenever the Amazon EC2 CreateImage API operation is called within the company's account.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS Lambda function to query AWS CloudTrail logs and to send an alert when a CreateImage API call is detected.
B. Configure AWS CloudTrail with an Amazon Simple Notification Service (Amazon SNS) notification that occurs when updated logs are sent to Amazon S3. Use Amazon Athena to create a new table and to query on CreateImage when an API call is detected.
C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for the CreateImage API cal
D. Configure the target as an Amazon Simple Notification Service (Amazon SNS) topic to send an alert when a CreateImage API call is detected.
E. Configure an Amazon Simple Queue Service (Amazon SQS) FIFO queue as a target for AWS CloudTrail log
F. Create an AWS Lambda function to send an alert to an Amazon Simple Notification Service (Amazon SNS) topic when a CreateImage API call is detected.

**Answer:** C

**Explanation:**
https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/monitor-ami-events.html#:~:text=For%20example%2C%20you%20can%20create%20an%20EventBridge%20rule%20that%20detects%20when%20the%20AMI%20creation%20process%20has%20completed%20and%20then%20invokes%20an%20Amazon%20SNS%20topic%20to% 20send%20an%20email%20notification%20to%20you.

**NEW QUESTION 91**
- (Topic 4)
A company runs a web application on Amazon EC2 instances in an Auto Scaling group that has a target group. The company desgned the application to work with session affinity (sticky sessions) for a better user experience.
The application must be available publicly over the internet as an endpoint_ A WAF must be applied to the endpoint for additional security. Session affinity (sticky sessions) must be configured on the endpoint
Which combination of steps will meet these requirements? (Select TWO)

A. Create a public Network Load Balancer Specify the application target group.
B. Create a Gateway Load Balancer Specify the application target group.
C. Create a public Application Load Balancer Specify the application target group.
D. Create a second target grou
E. Add Elastic IP addresses to the EC2 instances
F. Create a web ACL in AWS WAF Associate the web ACL with the endpoint

**Answer:** CE

**Explanation:**
C and E are the correct answers because they allow the company to create a public endpoint for its web application that supports session affinity (sticky sessions) and has a WAF applied for additional security. By creating a public Application Load Balancer, the company can distribute incoming traffic across multiple EC2 instances in an Auto Scaling group and specify the application target group. By creating a web ACL in AWS WAF and associating it with the Application Load Balancer, the company can protect its web application from common web exploits. By enabling session stickiness on the

Application Load Balancer, the company can ensure that subsequent requests from a user during a session are routed to the same target. References:
? Application Load Balancers
? AWS WAF
? Target Groups for Your Application Load Balancers
? How Application Load Balancer Works with Sticky Sessions

**NEW QUESTION 94**
- (Topic 4)
A company needs to store data from its healthcare application. The application's data frequently changes. A new regulation requires audit z access at all levels of the stored data.
The company hosts the application on an on-premises infrastructure that is running out of storage capacity. A solutions architect must securely migrate the existing data to AWS while satisfying the new regulation.
Which solution will meet these requirements?

A. Use AWS DataSync to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
B. Use AWS Snowcone to move the existing data to Amazon $3. Use AWS CloudTrail to log management events.
C. Use Amazon S3 Transfer Acceleration to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
D. Use AWS Storage Gateway to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.

**Answer:** A

**Explanation:**
This answer is correct because it meets the requirements of securely migrating the existing data to AWS and satisfying the new regulation. AWS DataSync is a service that makes it easy to move large amounts of data online between on-premises storage and Amazon S3. DataSync automatically encrypts data in transit and verifies data integrity during transfer. AWS CloudTrail is a service that records AWS API calls for your account and delivers log files to Amazon S3. CloudTrail can log data events, which show the resource operations performed on or within a resource in your AWS account, such as S3 object-level API activity. By using CloudTrail to log data events, you can audit access at all levels of the stored data.
References:
? https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html
? https://docs.aws.amazon.com/awscloudtrail/latest/userguide/logging-data-events- with-cloudtrail.html

**NEW QUESTION 99**
- (Topic 4)
A global marketing company has applications that run in the ap-southeast-2 Region and the eu-west-1 Region. Applications that run in a VPC in eu-west-1 need to communicate securely with databases that run in a VPC in ap-southeast-2.
Which network design will meet these requirements?

A. Create a VPC peering connection between the eu-west-1 VPC and the ap-southeast-2 VP
B. Create an inbound rule in the eu-west-1 application security group that allows traffic from the database server IP addresses in the ap-southeast-2 security group.
C. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west- 1 VP
D. Update the subnet route table
E. Create an inbound rule in the ap-southeast-2 database security group that references the security group ID of the application servers in eu-west-1.
F. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west- 1 VP
G. Update the subnet route tables Create an inbound rule in the ap-southeast-2 database security group that allows traffic from the eu-west-1 application server IP addresses.
H. Create a transit gateway with a peering attachment between the eu-west-1 VPC and the ap-southeast-2 VP
I. After the transit gateways are properly peered and routing is configured, create an inbound rule in the database security group that references the security group ID of the application servers in eu-west-1.

**Answer:** C

**Explanation:**
"You cannot reference the security group of a peer VPC that's in a different Region. Instead, use the CIDR block of the peer VPC."
https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html

**NEW QUESTION 101**
- (Topic 4)
A company has a large workload that runs every Friday evening. The workload runs on Amazon EC2 instances that are in two Availability Zones in the us-east-1 Region. Normally, the company must run no more than two instances at all times. However, the company wants to scale up to six instances each Friday to handle a regularly repeating increased workload.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a reminder in Amazon EventBridge to scale the instances.
B. Create an Auto Scaling group that has a scheduled action.
C. Create an Auto Scaling group that uses manual scaling.
D. Create an Auto Scaling group that uses automatic scaling.

**Answer:** B

**Explanation:**
An Auto Scaling group is a collection of EC2 instances that share similar characteristics and can be scaled in or out automatically based on demand. An Auto Scaling group can have a scheduled action, which is a configuration that tells the group to scale to a specific size at a specific time. This way, the company can scale up to six instances each Friday evening to handle the increased workload, and scale down to two instances at other times to save costs. This solution meets the requirements with the least operational overhead, as it does not require manual intervention or custom scripts. References:
? 1 explains how to create a scheduled action for an Auto Scaling group.
? 2 describes the concept and benefits of an Auto Scaling group.

**NEW QUESTION 104**

- (Topic 4)
A company has an online gaming application that has TCP and UDP multiplayer gaming capabilities. The company uses Amazon Route 53 to point the application traffic to multiple Network Load Balancers (NLBs) in different AWS Regions. The company needs to improve application performance and decrease latency for the online game in preparation for user growth.
Which solution will meet these requirements?

A. Add an Amazon CloudFront distribution in front of the NLB
B. Increase the Cache- Control: max-age parameter.
C. Replace the NLBs with Application Load Balancers (ALBs). Configure Route 53 to use latency-based routing.
D. Add AWS Global Accelerator in front of the NLB
E. Configure a Global Accelerator endpoint to use the correct listener ports.
F. 'Add an Amazon API Gateway endpoint behind the NLB
G. Enable API cachin
H. Override method caching for the different stages.

**Answer:** C

**Explanation:**
This answer is correct because it improves the application performance and
decreases latency for the online game by using AWS Global Accelerator. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your public applications. Global Accelerator provides two global static public IPs that act as a fixed entry point to your application endpoints, such as NLBs, in different AWS Regions. Global Accelerator uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure. Global Accelerator also terminates TCP and UDP traffic at the edge locations, which reduces the number of hops and improves the network performance. By adding AWS Global Accelerator in front of the NLBs, you can achieve up to 60% improvement in latency for your online game.
References:
? https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global- accelerator.html
? https://aws.amazon.com/global-accelerator/

**NEW QUESTION 106**
- (Topic 4)
An loT company is releasing a mattress that has sensors to collect data about a user's sleep. The sensors will send data to an Amazon S3 bucket. The sensors collect approximately 2 MB of data every night for each mattress. The company must process and summarize the data for each mattress. The results need to be available as soon as possible Data processing will require 1 GB of memory and will finish within 30 seconds.
Which solution will meet these requirements MOST cost-effectively?

A. Use AWS Glue with a Scalajob.
B. Use Amazon EMR with an Apache Spark script.
C. Use AWS Lambda with a Python script.
D. Use AWS Glue with a PySpark job.

**Answer:** C

**Explanation:**
AWS Lambda charges you based on the number of invocations and the execution time of your function. Since the data processing job is relatively small (2 MB of data), Lambda is a cost-effective choice. You only pay for the actual usage without the need to provision and maintain infrastructure.

**NEW QUESTION 109**
- (Topic 4)
A company is moving its on-premises Oracle database to Amazon Aurora PostgreSQL. The database has several applications that write to the same tables. The applications need to be migrated one by one with a month in between each migration. Management has expressed concerns that the database has a high number of reads and writes. The data must be kept in sync across both databases throughout the migration.
What should a solutions architect recommend?

A. Use AWS DataSync for the initial migratio
B. Use AWS Database Migration Service (AWS DMS) to create a change data capture (CDC) replication task and a table mapping to select all tables.
C. Use AWS DataSync for the initial migratio
D. Use AWS Database Migration Service (AWS DMS) to create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
E. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a memory optimized replication instanc
F. Create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
G. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a compute optimized replication instanc
H. Create a full load plus change data capture (CDC) replication task and a table mapping to select the largest tables.

**Answer:** C

**Explanation:**
https://aws.amazon.com/ko/premiumsupport/knowledge-center/dms- memory-optimization/

**NEW QUESTION 111**
- (Topic 4)
A company has NFS servers in an on-premises data center that need to periodically back up small amounts of data to Amazon S3. Which solution meets these requirements and is MOST cost-effective?

A. Set up AWS Glue to copy the data from the on-premises servers to Amazon S3.
B. Set up an AWS DataSync agent on the on-premises servers, and sync the data to Amazon S3.
C. Set up an SFTP sync using AWS Transfer for SFTP to sync data from on premises to Amazon S3.
D. Set up an AWS Direct Connect connection between the on-premises data center and a VPC, and copy the data to Amazon S3.

**Answer:** B

**Explanation:**
AWS DataSync is a service that makes it easy to move large amounts of data online between on-premises storage and AWS storage services. AWS DataSync can transfer data at speeds up to 10 times faster than open-source tools by using a purpose- built network protocol and parallelizing data transfers. AWS DataSync also handles encryption, data integrity verification, and bandwidth optimization. To use AWS DataSync, users need to deploy a DataSync agent on their on-premises servers, which connects to the NFS servers and syncs the data to Amazon S3. Users can schedule periodic or one- time sync tasks and monitor the progress and status of the transfers.
The other options are not correct because they are either not cost-effective or not suitable for the use case. Setting up AWS Glue to copy the data from the on-premises servers to Amazon S3 is not cost-effective because AWS Glue is a serverless data integration service that is mainly used for extract, transform, and load (ETL) operations, not for simple data backup. Setting up an SFTP sync using AWS Transfer for SFTP to sync data from on premises to Amazon S3 is not cost-effective because AWS Transfer for SFTP is a fully managed service that provides secure file transfer using the SFTP protocol, which is more suitable for exchanging data with third parties than for backing up data. Setting up an AWS Direct Connect connection between the on-premises data center and a VPC, and copying the data to Amazon S3 is not cost-effective because AWS Direct Connect is a dedicated network connection between AWS and the on-premises location, which has high upfront costs and requires additional configuration.
References:
? AWS DataSync
? How AWS DataSync works
? AWS DataSync FAQs

**NEW QUESTION 113**
- (Topic 4)
A manufacturing company has machine sensors that upload .csv files to an Amazon S3 bucket. These .csv files must be converted into images and must be made available as soon as possible for the automatic generation of graphical reports.
The images become irrelevant after 1 month, but the .csv files must be kept to train machine learning (ML) models twice a year. The ML trainings and audits are planned weeks in advance.
Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

A. Launch an Amazon EC2 Spot Instance that downloads the .csv files every hour, generates the image files, and uploads the images to the S3 bucket.
B. Design an AWS Lambda function that converts the .csv files into images and stores the images in the S3 bucke
C. Invoke the Lambda function when a .csv file is uploaded.
D. Create S3 Lifecycle rules for .csv files and image files in the S3 bucke
E. Transition the.csv files from S3 Standard to S3 Glacier 1 day after they are uploade
F. Expire the image files after 30 days.
G. Create S3 Lifecycle rules for .csv files and image files in the S3 bucke
H. Transition the.csv files from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) 1 day after they are uploade
I. Expire the image files after 30 days.
J. Create S3 Lifecycle rules for .csv files and image files in the S3 bucke
K. Transition the.csv files from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) 1 day after they are uploade
L. Keep the image files in Reduced Redundancy Storage (RRS).

**Answer:** BC

**Explanation:**
These answers are correct because they meet the requirements of
converting the .csv files into images, making them available as soon as possible, and minimizing the storage costs. AWS Lambda is a service that lets you run code without provisioning or managing servers. You can use AWS Lambda to design a function that converts the .csv files into images and stores the images in the S3 bucket. You can invoke the Lambda function when a .csv file is uploaded to the S3 bucket by using an S3 event notification. This way, you can ensure that the images are generated and made available as soon as possible for the graphical reports. S3 Lifecycle is a feature that enables you to manage your objects so that they are stored cost effectively throughout their lifecycle. You can create S3 Lifecycle rules for .csv files and image files in the S3 bucket to transition them to different storage classes or expire them based on your business needs. You can transition the .csv files from S3 Standard to S3 Glacier 1 day after they are uploaded, since they are only needed twice a year for ML trainings and audits that are planned weeks in advance. S3 Glacier is a storage class for data archiving that offers secure, durable, and extremely low-cost storage with retrieval times ranging from minutes to hours. You can expire the image files after 30 days, since they become irrelevant after 1 month. References:
? https://docs.aws.amazon.com/lambda/latest/dg/welcome.html
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/NotificationHowTo.html
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle- mgmt.html
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class- intro.html#sc-glacier

**NEW QUESTION 117**
- (Topic 4)
A company has applications hosted on Amazon EC2 instances with IPv6 addresses. The applications must initiate communications with other external applications using the internet.
However, the company’s security policy states that any external service cannot initiate a connection to the EC2 instances.
What should a solutions architect recommend to resolve this issue?

A. Create a NAT gateway and make it the destination of the subnet's route table.
B. Create an internet gateway and make it the destination of the subnet's route table
C. Create a virtual private gateway and make it the destination of the subnet's route table.
D. Create an egress-only internet gateway and make it the destination of the subnet's route table.

**Answer:** D

**Explanation:**
An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in your VPC to the internet, and prevents the internet from initiating an IPv6 connection with your instances. This meets the company’s security policy and requirements. To use an egress-only internet gateway, you need to add a route in the subnet’s route table that routes IPv6 internet traffic (::/0) to the egress-only internet gateway.
Reference URLs:
1 https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html
2 https://dev.to/aws-builders/what-is-an-egress-only-internet-gateways-in-aws-7gp
3 https://docs.aws.amazon.com/vpc/latest/userguide/route-table-options.html

**NEW QUESTION 119**
- (Topic 4)
An ecommerce company runs applications in AWS accounts that are part of an organization in AWS Organizations The applications run on Amazon Aurora PostgreSQL databases across all the accounts The company needs to prevent malicious activity and must identify abnormal failed and incomplete login attempts to the databases
Which solution will meet these requirements in the MOST operationally efficient way?

A. Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts
B. Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization
C. Publish the Aurora general logs to a log group in Amazon CloudWatch Logs Export the log data to a central Amazon S3 bucket
D. Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket

**Answer:** C

**Explanation:**
This option is the most operationally efficient way to meet the requirements because it allows the company to monitor and analyze the database login activity across all the accounts in the organization. By publishing the Aurora general logs to a log group in Amazon CloudWatch Logs, the company can enable the logging of the database connections, disconnections, and failed authentication attempts. By exporting the log data to a central Amazon S3 bucket, the company can store the log data in a durable and cost- effective way and use other AWS services or tools to perform further analysis or alerting on the log data. For example, the company can use Amazon Athena to query the log data in Amazon S3, or use Amazon SNS to send notifications based on the log data.
* A. Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts. This option is not effective because SCPs are not designed to identify the failed login attempts, but to restrict the actions that the users and roles can perform in the member accounts of the organization. SCPs are applied to the AWS API calls, not to the database login attempts. Moreover, SCPs do not provide any logging or analysis capabilities for the database activity.
* B. Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization. This option is not optimal because the Amazon RDS Protection feature in Amazon GuardDuty is not available for Aurora PostgreSQL databases, but only for Amazon RDS for MySQL and Amazon RDS for MariaDB databases. Moreover, the Amazon RDS Protection feature does not monitor the database login attempts, but the network and API activity related to the RDS instances.
* D. Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket. This option is not sufficient because AWS CloudTrail does not capture the database login attempts, but only the AWS API calls made by or on behalf of the Aurora PostgreSQL database. For example, AWS CloudTrail can record the events such as creating, modifying, or deleting the database instances, clusters, or snapshots, but not the events such as connecting, disconnecting, or failing to authenticate to the database. References:
? 1 Working with Amazon Aurora PostgreSQL - Amazon Aurora
? 2 Working with log groups and log streams - Amazon CloudWatch Logs
? 3 Exporting Log Data to Amazon S3 - Amazon CloudWatch Logs
? [4] Amazon GuardDuty FAQs
? [5] Logging Amazon RDS API Calls with AWS CloudTrail - Amazon Relational Database Service

**NEW QUESTION 124**
- (Topic 4)
A company has a nightly batch processing routine that analyzes report files that an on- premises file system receives daily through SFTP. The company wants to move the solution to the AWS Cloud. The solution must be highly available and resilient. The solution also must minimize operational effort.
Which solution meets these requirements?

A. Deploy AWS Transfer for SFTP and an Amazon Elastic File System (Amazon EFS) file system for storag
B. Use an Amazon EC2 instance in an Auto Scaling group with a scheduled scaling policy to run the batch operation.
C. Deploy an Amazon EC2 instance that runs Linux and an SFTP servic
D. Use an Amazon Elastic Block Store {Amazon EBS) volume for storag
E. Use an Auto Scaling group with the minimum number of instances and desired number of instances set to 1.
F. Deploy an Amazon EC2 instance that runs Linux and an SFTP servic
G. Use an Amazon Elastic File System (Amazon EFS) file system for storag
H. Use an Auto Scaling group with the minimum number of instances and desired number of instances set to 1.
I. Deploy AWS Transfer for SFTP and an Amazon S3 bucket for storag
J. Modify the application to pull the batch files from Amazon S3 to an Amazon EC2 instance for processin
K. Use an EC2 instance in an Auto Scaling group with a scheduled scaling policy to run the batch operation.

**Answer:** D

**Explanation:**
The solution that meets the requirements of high availability, performance, security, and static IP addresses is to use Amazon CloudFront, Application Load Balancers (ALBs), Amazon Route 53, and AWS WAF. This solution allows the company to distribute its HTTP-based application globally using CloudFront, which is a content delivery network (CDN) service that caches content at edge locations and provides static IP addresses for each edge location. The company can also use Route 53 latency-based routing to route requests to the closest ALB in each Region, which balances the load across the EC2 instances. The company can also deploy AWS WAF on the CloudFront distribution to protect the application against common web exploits by creating rules that allow, block, or count web requests based on conditions that are defined. The other solutions do not meet all the requirements because they either use Network Load Balancers (NLBs), which do not support HTTP-based applications, or they do not use CloudFront, which provides better performance and security than AWS Global Accelerator.
References :=
? Amazon CloudFront
? Application Load Balancer
? Amazon Route 53
? AWS WAF

**NEW QUESTION 128**
- (Topic 4)
A company uses high concurrency AWS Lambda functions to process a constantly increasing number of messages in a message queue during marketing events. The Lambda functions use CPU intensive code to process the messages. The company wants to reduce the compute costs and to maintain service latency for its customers.
Which solution will meet these requirements?

A. Configure reserved concurrency for the Lambda function
B. Decrease the memory allocated to the Lambda functions.
C. Configure reserved concurrency for the Lambda function

D. Increase the memory according to AWS Compute Optimizer recommendations.
E. Configure provisioned concurrency for the Lambda function
F. Decrease the memory allocated to the Lambda functions.
G. Configure provisioned concurrency for the Lambda function
H. Increase the memory according to AWS Compute Optimizer recommendations.

**Answer:** D

**Explanation:**
The company wants to reduce the compute costs and maintain service latency for its Lambda functions that process a constantly increasing number of messages in a message queue. The Lambda functions use CPU intensive code to process the messages. To meet these requirements, a solutions architect should recommend the following solution:
? Configure provisioned concurrency for the Lambda functions. Provisioned
concurrency is the number of pre-initialized execution environments that are allocated to the Lambda functions. These execution environments are prepared to respond immediately to incoming function requests, reducing the cold start latency. Configuring provisioned concurrency also helps to avoid throttling errors due to reaching the concurrency limit of the Lambda service.
? Increase the memory according to AWS Compute Optimizer recommendations.
AWS Compute Optimizer is a service that provides recommendations for optimal AWS resource configurations based on your utilization data. By increasing the memory allocated to the Lambda functions, you can also increase the CPU power and improve the performance of your CPU intensive code. AWS Compute Optimizer can help you find the optimal memory size for your Lambda functions based on your workload characteristics and performance goals.
This solution will reduce the compute costs by avoiding unnecessary over-provisioning of memory and CPU resources, and maintain service latency by using provisioned concurrency and optimal memory size for the Lambda functions.
References:
? Provisioned Concurrency
? AWS Compute Optimizer

**NEW QUESTION 131**
- (Topic 4)
A company uses multiple vendors to distribute digital assets that are stored in Amazon S3
buckets The company wants to ensure that its vendor AWS accounts have the minimum access that is needed to download objects in these S3 buckets
Which solution will meet these requirements with the LEAST operational overhead?

A. Design a bucket policy that has anonymous read permissions and permissions to list ail buckets.
B. Design a bucket policy that gives read-only access to user
C. Specify 1AM entities as principals
D. Create a cross-account 1AM role that has a read-only access policy specified for the 1AM role.
E. Create a user policy and vendor user groups that give read-only access to vendor users

**Answer:** C

**Explanation:**
A cross-account IAM role is a way to grant users from one AWS account access to resources in another AWS account. The cross-account IAM role can have a read-only access policy attached to it, which allows the users to download objects from the S3 buckets without modifying or deleting them. The cross-account IAM role also reduces the operational overhead of managing multiple IAM users and policies in each account. The cross-account IAM role meets all the requirements of the question, while the other options do not. References:
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-walkthroughs-
managing-access-example2.html
? https://aws.amazon.com/blogs/storage/setting-up-cross-account-amazon-s3- access-with-s3-access-points/
? https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for- user_externalid.html

**NEW QUESTION 132**
- (Topic 4)
A company wants to run its experimental workloads in the AWS Cloud. The company has a budget for cloud spending. The company's CFO is concerned about cloud spending accountability for each department. The CFO wants to receive notification when the spending threshold reaches 60% of the budget.
Which solution will meet these requirements?

A. Use cost allocation tags on AWS resources to label owner
B. Create usage budgets in AWS Budget
C. Add an alert threshold to receive notification when spending exceeds 60% of the budget.
D. Use AWS Cost Explorer forecasts to determine resource owner
E. Use AWS Cost Anomaly Detection to create alert threshold notifications when spending exceeds 60% of the budget.
F. Use cost allocation tags on AWS resources to label owner
G. Use AWS Support API on AWS Trusted Advisor to create alert threshold notifications when spending exceeds 60% of the budget
H. Use AWS Cost Explorer forecasts to determine resource owner
I. Create usage budgets in AWS Budget
J. Add an alert threshold to receive notification when spending exceeds 60% of the budget.

**Answer:** A

**Explanation:**
This solution meets the requirements because it allows the company to track and manage its cloud spending by using cost allocation tags to assign costs to different departments, creating usage budgets to set spending limits, and adding alert thresholds to receive notifications when the spending reaches a certain percentage of the budget. This way, the company can monitor its experimental workloads and avoid overspending on the cloud.
References:
? Using Cost Allocation Tags
? Creating an AWS Budget
? Creating an Alert for an AWS Budget

**NEW QUESTION 134**
- (Topic 4)

A gaming company uses Amazon DynamoDB to store user information such as geographic location, player data, and leaderboards. The company needs to configure continuous backups to an Amazon S3 bucket with a minimal amount of coding. The backups must not affect availability of the application and must not affect the read capacity units (RCUs) that are defined for the table
Which solution meets these requirements?

A. Use an Amazon EMR cluste
B. Create an Apache Hive job to back up the data to Amazon S3.
C. Export the data directly from DynamoDB to Amazon S3 with continuous backup
D. Turn on point-in-time recovery for the table.
E. Configure Amazon DynamoDB Stream
F. Create an AWS Lambda function to consume the stream and export the data to an Amazon S3 bucket.
G. Create an AWS Lambda function to export the data from the database tables to Amazon S3 on a regular basi
H. Turn on point-in-time recovery for the table.

**Answer:** B

**Explanation:**
https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and- design-patterns/
https://aws.amazon.com/premiumsupport/knowledge-center/back-up-dynamodb-s3/

**NEW QUESTION 135**
- (Topic 4)
A company's web application that is hosted in the AWS Cloud recently increased in popularity. The web application currently exists on a single Amazon EC2 instance in a single public subnet. The web application has not been able to meet the demand of the increased web traffic.
The company needs a solution that will provide high availability and scalability to meet the increased user demand without rewriting the web application.
Which combination of steps will meet these requirements? (Select TWO.)

A. Replace the EC2 instance with a larger compute optimized instance.
B. Configure Amazon EC2 Auto Scaling with multiple Availability Zones in private subnets.
C. Configure a NAT gateway in a public subnet to handle web requests.
D. Replace the EC2 instance with a larger memory optimized instance.
E. Configure an Application Load Balancer in a public subnet to distribute web traffic

**Answer:** BE

**Explanation:**
These two steps will meet the requirements because they will provide high availability and scalability for the web application without rewriting it. Amazon EC2 Auto Scaling allows you to automatically adjust the number of EC2 instances in response to changes in demand. By configuring Auto Scaling with multiple Availability Zones in private subnets, you can ensure that your web application is distributed across isolated and fault-tolerant locations, and that your instances are not directly exposed to the internet. An Application Load Balancer operates at the application layer and distributes incoming web traffic across multiple targets, such as EC2 instances, containers, or Lambda functions. By configuring an Application Load Balancer in a public subnet, you can enable your web application to handle requests from the internet and route them to the appropriate targets in the private subnets.
References:
? What is Amazon EC2 Auto Scaling?
? What is an Application Load Balancer?

**NEW QUESTION 138**
- (Topic 4)
The following IAM policy is attached to an IAM group. This is the only policy applied to the group.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "1",
            "Effect": "Allow",
            "Action": "ec2:*",
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "ec2:Region": "us-east-1"
                }
            }
        },
        {
            "Sid": "2",
            "Effect": "Deny",
            "Action": [
                "ec2:StopInstances",
                "ec2:TerminateInstances"
            ],
            "Resource": "*",
            "Condition": {
                "BoolIfExists": {"aws:MultiFactorAuthPresent": false}
            }
        }
    ]
}
```

A. Group members are permitted any Amazon EC2 action within the us-east-1 Regio
B. Statements after the Allow permission are not applied.
C. Group members are denied any Amazon EC2 permissions in the us-east-1 Region unless they are logged in with multi-factor authentication (MFA).
D. Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for all Regions when logged in with multi-factor authentication (MFA). Group members arepermitted any other Amazon EC2 action.
E. Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for the us-east-1 Region only when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action within the us-east-1 Region.

**Answer:** D

**Explanation:**
This answer is correct because it reflects the effect of the IAM policy on the group members. The policy has two statements: one with an Allow effect and one with a Deny effect. The Allow statement grants permission to perform any EC2 action on any resource within the us-east-1 Region. The Deny statement overrides the Allow statement and denies permission to perform the ec2:StopInstances and ec2:TerminateInstances actions on any resource within the us-east-1 Region, unless the group member is logged in with MFA. Therefore, the group members can perform any EC2 action except stopping or terminating instances in the us-east-1 Region, unless they use MFA.

**NEW QUESTION 141**
- (Topic 4)
A security audit reveals that Amazon EC2 instances are not being patched regularly. A solutions architect needs to provide a solution that will run regular security scans across a
large fleet of EC2 instances. The solution should also patch the EC2 instances on a regular schedule and provide a report of each instance's patch status.
Which solution will meet these requirements?

A. Set up Amazon Macie to scan the EC2 instances for software vulnerabilitie
B. Set up a cron job on each EC2 instance to patch the instance on a regular schedule.
C. Turn on Amazon GuardDuty in the accoun
D. Configure GuardDuty to scan the EC2 instances for software vulnerabilitie
E. Set up AWS Systems Manager Session Manager to patch the EC2 instances on a regular schedule.
F. Set up Amazon Detective to scan the EC2 instances for software vulnerabilitie
G. Set up an Amazon EventBridge scheduled rule to patch the EC2 instances on a regular schedule.
H. Turn on Amazon Inspector in the accoun
I. Configure Amazon Inspector to scan the EC2 instances for software vulnerabilitie
J. Set up AWS Systems Manager Patch Manager to patch the EC2 instances on a regular schedule.

**Answer:** D

**Explanation:**
Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for exposure, vulnerabilities, and deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity1. Amazon Inspector can scan the EC2 instances for software vulnerabilities and provide a report of each instance’s patch status. AWS Systems Manager Patch Manager is a capability of AWS Systems Manager that automates the process of patching managed nodes with both security-related updates and other types of updates. Patch Manager uses patch baselines, which include rules for auto- approving patches within days of their release, in addition to optional lists of approved and rejected patches. Patch Manager can patch fleets of Amazon EC2 instances, edge devices, on-premises servers, and virtual machines (VMs) by operating system type2. Patch Manager can patch the EC2 instances on a regular schedule and provide a report of each instance’s patch status. Therefore, the combination of Amazon Inspector and AWS Systems Manager Patch Manager will meet the

requirements of the question.
The other options are not valid because:
? Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie does not scan the EC2 instances for software vulnerabilities, but rather for data classification and protection3. A cron job is a Linux command for scheduling a task to be executed sometime in the future. A cron job is not a reliable way to patch the EC2 instances on a regular schedule, as it may fail or be interrupted by other processes4.
? Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts and workloads. Amazon GuardDuty does not scan the EC2 instances for software vulnerabilities, but rather for network and API activity anomalies5. AWS Systems Manager Session Manager is a fully managed AWS Systems Manager capability that lets you manage your Amazon EC2 instances, edge devices, on-premises servers, and virtual machines (VMs) through an interactive one-click browser-based shell or the AWS Command Line Interface (AWS CLI). Session Manager does not patch the EC2 instances on a regular schedule, but rather provides secure and auditable node management2.
? Amazon Detective is a security service that makes it easy to analyze, investigate, and quickly identify the root cause of potential security issues or suspicious activities. Amazon Detective does not scan the EC2 instances for software vulnerabilities, but rather collects and analyzes data from AWS sources such as Amazon GuardDuty, Amazon VPC Flow Logs, and AWS CloudTrail. Amazon EventBridge is a serverless event bus that makes it easy to connect applications using data from your own applications, integrated Software-as-a-Service (SaaS) applications, and AWS services. EventBridge delivers a stream of real-time data from event sources, such as Zendesk, Datadog, or Pagerduty, and routes that data to targets like AWS Lambda. EventBridge does not patch the EC2 instances on a regular schedule, but rather triggers actions based on events.
References: Amazon Inspector, AWS Systems Manager Patch Manager, Amazon Macie, Cron job, Amazon GuardDuty, [Amazon Detective], [Amazon EventBridge]

**NEW QUESTION 143**
- (Topic 4)
A company runs a highly available web application on Amazon EC2 instances behind an Application Load Balancer The company uses Amazon CloudWatch metrics
As the traffic to the web application Increases, some EC2 instances become overloaded with many outstanding requests The CloudWatch metrics show that the number of requests processed and the time to receive the responses from some EC2 instances are both higher compared to other EC2 instances The company does not want new requests to be forwarded to the EC2 instances that are already overloaded.
Which solution will meet these requirements?

A. Use the round robin routing algorithm based on the RequestCountPerTarget and Active Connection Count CloudWatch metrics.
B. Use the least outstanding requests algorithm based on the RequestCountPerTarget and ActiveConnectionCount CloudWatch metrics.
C. Use the round robin routing algorithm based on the RequestCount and TargetResponseTime CloudWatch metrics.
D. Use the least outstanding requests algorithm based on the RequestCount and TargetResponseTime CloudWatch metrics.

**Answer:** D

**Explanation:**
The least outstanding requests (LOR) algorithm is a load balancing algorithm that distributes incoming requests to the target with the fewest outstanding requests. This helps to avoid overloading any single target and improves the overall performance and availability of the web application. The LOR algorithm can use the RequestCount and TargetResponseTime CloudWatch metrics to determine the number of outstanding requests and the response time of each target. These metrics measure the number of requests processed by each target and the time elapsed after the request leaves the load balancer until a response from the target is received by the load balancer, respectively. By using these metrics, the LOR algorithm can route new requests to the targets that are less busy and more responsive, and avoid sending requests to the targets that are already overloaded or slow. This solution meets the requirements of the company.
References:
? Application Load Balancer now supports Least Outstanding Requests algorithm for load balancing requests
? Target groups for your Application Load Balancers
? Elastic Load Balancing - Application Load Balancers

**NEW QUESTION 148**
- (Topic 4)
A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should.be protected throughout the entire application stack, and access to the information should be restricted to certain applications.
Which action should the solutions architect take?

A. Configure a CloudFront signed URL.
B. Configure a CloudFront signed cookie.
C. Configure a CloudFront field-level encryption profile.
D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

**Answer:** C

**Explanation:**
it allows the company to protect sensitive information submitted by users throughout the entire application stack and restrict access to certain applications. By configuring a CloudFront field-level encryption profile, the company can encrypt specific fields of user data at the edge locations before sending it to the origin servers. By using public-private key pairs, the company can ensure that only authorized applications can decrypt and access the sensitive information. References:
? Field-Level Encryption
? Encrypting and Decrypting Data

**NEW QUESTION 152**
- (Topic 4)
A company hosts multiple applications on AWS for different product lines. The applications use different compute resources, including Amazon EC2 instances and Application Load Balancers. The applications run in different AWS accounts under the same organization in AWS Organizations across multiple AWS Regions. Teams for each product line have tagged each compute resource in the individual accounts.
The company wants more details about the cost for each product line from the consolidated billing feature in Organizations.
Which combination of steps will meet these requirements? (Select TWO.)

A. Select a specific AWS generated tag in the AWS Billing console.

B. Select a specific user-defined tag in the AWS Billing console.
C. Select a specific user-defined tag in the AWS Resource Groups console.
D. Activate the selected tag from each AWS account.
E. Activate the selected tag from the Organizations management account.

**Answer:** BE

**Explanation:**
User-defined tags are key-value pairs that can be applied to AWS resources to categorize and track them. User-defined tags can also be used to allocate costs and create detailed billing reports in the AWS Billing console. To use user-defined tags for cost allocation, the tags must be activated from the Organizations management account, which is the root account that has full control over all the member accounts in the organization. Once activated, the user-defined tags will appear as columns in the cost allocation report, and can be used to filter and group costs by product line. This solution will meet the requirements with the least operational overhead, as it leverages the existing tagging strategy and does not require any code development or manual intervention.
References:
? 1 explains how to use user-defined tags for cost allocation.
? 2 describes how to access and manage member accounts from the Organizations management account.
? 3 discusses how to create and view cost allocation reports in the AWS Billing console.

**NEW QUESTION 154**
- (Topic 4)
A company has a financial application that produces reports. The reports average 50 KB in size and are stored in Amazon S3. The reports are frequently accessed during the first week after production and must be stored for several years. The reports must be retrievable within 6 hours.
Which solution meets these requirements MOST cost-effectively?

A. Use S3 Standar
B. Use an S3 Lifecycle rule to transition the reports to S3 Glacier after 7 days.
C. Use S3 Standar
D. Use an S3 Lifecycle rule to transition the reports to S3 Standard- Infrequent Access (S3 Standard-IA) after 7 days.
E. Use S3 Intelligent-Tierin
F. Configure S3 Intelligent-Tiering to transition the reports to S3 Standard-Infrequent Access (S3 Standard-IA) and S3 Glacier.
G. Use S3 Standar
H. Use an S3 Lifecycle rule to transition the reports to S3 Glacier Deep Archive after 7 days.

**Answer:** A

**Explanation:**
To store and retrieve reports that are frequently accessed during the first week and must be stored for several years, S3 Standard and S3 Glacier are suitable solutions. S3 Standard offers high durability, availability, and performance for frequently accessed data. S3 Glacier offers secure and durable storage for long-term data archiving at a low cost. S3 Lifecycle rules can be used to transition the reports from S3 Standard to S3 Glacier after 7 days, which can reduce storage costs. S3 Glacier also supports retrieval within 6 hours.
References:
? Storage Classes
? Object Lifecycle Management
? Retrieving Archived Objects from Amazon S3 Glacier

**NEW QUESTION 155**
- (Topic 4)
A company wants to migrate its on-premises Microsoft SQL Server Enterprise edition database to AWS. The company's online application uses the database to process transactions. The data analysis team uses the same production database to run reports for analytical processing. The company wants to reduce operational overhead by moving to managed services wherever possible.
Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate to Amazon RDS for Microsoft SQL Serve
B. Use read replicas for reporting purposes.
C. Migrate to Microsoft SQL Server on Amazon EC2. Use Always On read replicas for reporting purposes.
D. Migrate to Amazon DynamoD
E. Use DynamoDB on-demand replicas for reporting purposes.
F. Migrate to Amazon Aurora MySQ
G. Use Aurora read replicas for reporting purposes.

**Answer:** A

**Explanation:**
Amazon RDS for Microsoft SQL Server is a fully managed service that offers SQL Server 2014, 2016, 2017, and 2019 editions while offloading database administration tasks such as backups, patching, and scaling. Amazon RDS supports read replicas, which are read-only copies of the primary database that can be used for reporting purposes without affecting the performance of the online application. This solution will meet the requirements with the least operational overhead, as it does not require any code changes or manual intervention.
References:
? 1 provides an overview of Amazon RDS for Microsoft SQL Server and its benefits.
? 2 explains how to create and use read replicas with Amazon RDS.

**NEW QUESTION 158**
- (Topic 4)
A company runs an application on a group of Amazon Linux EC2 instances. For compliance reasons, the company must retain all application log files for 7 years. The log files will be analyzed by a reporting tool that must be able to access all the files concurrently.
Which storage solution meets these requirements MOST cost-effectively?

A. Amazon Elastic Block Store (Amazon EBS)
B. Amazon Elastic File System (Amazon EFS)
C. Amazon EC2 instance store

D. Amazon S3

**Answer:** D

**Explanation:**
https://docs.aws.amazon.com/efs/latest/ug/transfer-data-to-efs.html

**NEW QUESTION 160**
- (Topic 4)
A company is deploying a new application on Amazon EC2 instances. The application writes data to Amazon Elastic Block Store (Amazon EBS) volumes. The company needs to ensure that all data that is written to the EBS volumes is encrypted at rest.
Which solution will meet this requirement?

A. Create an 1AM role that specifies EBS encryptio
B. Attach the role to the EC2 instances.
C. Create the EBS volumes as encrypted volume
D. Attach the EBS volumes to the EC2 instances
E. Create an EC2 instance tag that has a key of Encrypt and a value of Tru
F. Tag all instances that require encryption at the EBS level.
G. Create an AWS Key Management Service (AWS KMS) key policy that enforces EBSencryption in the accoun
H. Ensure that the key policy is active

**Answer:** B

**Explanation:**
The solution that will meet the requirement of ensuring that all data that is written to the EBS volumes is encrypted at rest is B. Create the EBS volumes as encrypted volumes and attach the encrypted EBS volumes to the EC2 instances. When you create an EBS volume, you can specify whether to encrypt the volume. If you choose to encrypt the volume, all data written to the volume is automatically encrypted at rest using AWS- managed keys. You can also use customer-managed keys (CMKs) stored in AWS KMS to encrypt and protect your EBS volumes. You can create encrypted EBS volumes and attach them to EC2 instances to ensure that all data written to the volumes is encrypted at rest.

**NEW QUESTION 162**
.....

# Thank You for Trying Our Product

## We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

## AWS-Solution-Architect-Associate Practice Exam Features:

* AWS-Solution-Architect-Associate Questions and Answers Updated Frequently

* AWS-Solution-Architect-Associate Practice Questions Verified by Expert Senior Certified Staff

* AWS-Solution-Architect-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The AWS-Solution-Architect-Associate Practice Test Here