Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

NEW QUESTION 1
- (Topic 4)
A company has a business-critical application that runs on Amazon EC2 instances. The application stores data in an Amazon DynamoDB table. The company must be able to revert the table to any point within the last 24 hours.
Which solution meets these requirements with the LEAST operational overhead?

A. Configure point-in-time recovery for the table.
B. Use AWS Backup for the table.
C. Use an AWS Lambda function to make an on-demand backup of the table every hour.
D. Turn on streams on the table to capture a log of all changes to the table in the last 24 hours Store a copy of the stream in an Amazon S3 bucket.

Answer: A

Explanation:
Point-in-time recovery (PITR) for DynamoDB is a feature that enables you to restore your table data to any point in time during the last 35 days. PITR helps protect your table from accidental write or delete operations, such as a test script writing to a production table or a user issuing a wrong command. PITR is easy to use, fully managed, fast, and scalable. You can enable PITR with a single click in the DynamoDB console or with a simple API call. You can restore a table to a new table using the console, the AWS CLI, or the DynamoDB API. PITR does not consume any provisioned table capacity and has no impact on the performance or availability of your production applications. PITR meets the requirements of the company with the least operational overhead, as it does not require any manual backup creation, scheduling, or maintenance. It also provides per-second granularity for restoring the table to any point within the last 24 hours.
References:
? Point-in-time recovery for DynamoDB - Amazon DynamoDB
? Amazon DynamoDB point-in-time recovery (PITR)
? Enable Point-in-Time Recovery (PITR) for Dynamodb global tables
? Restoring a DynamoDB table to a point in time - Amazon DynamoDB
? Point-in-time recovery: How it works - Amazon DynamoDB

NEW QUESTION 2
- (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 3
- (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 4
- (Topic 4)
A company's ecommerce website has unpredictable traffic and uses AWS Lambda functions to directly access a private Amazon RDS for PostgreSQL DB instance. The company wants to maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections.
What should a solutions architect do to meet these requirements?

A. Point the client driver at an RDS custom endpoin
B. Deploy the Lambda functions inside a VPC.
C. Point the client driver at an RDS proxy endpoin
D. Deploy the Lambda functions inside a VPC.
E. Point the client driver at an RDS custom endpoin
F. Deploy the Lambda functions outside a VPC.
G. Point the client driver at an RDS proxy endpoin
H. Deploy the Lambda functions outside a VPC.

Answer: B

Explanation:
To maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections, a solutions architect should point the client driver at an RDS proxy endpoint and deploy the Lambda functions inside a VPC. An RDS proxy is a fully managed database proxy that allows applications to share connections to a database, improving database availability and scalability. By using an RDS proxy, the Lambda functions can reuse existing connections, rather than creating new ones for every invocation, reducing the connection overhead and latency. Deploying the Lambda functions inside a VPC allows them to access the private RDS DB instance securely and efficiently, without exposing it to the public internet. References:
? Using Amazon RDS Proxy with AWS Lambda
? Configuring a Lambda function to access resources in a VPC

NEW QUESTION 5
- (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 6
- (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 7
- (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 8
- (Topic 4)
A company needs to use its on-premises LDAP directory service to authenticate its users to the AWS Management Console. The directory service is not compatible with Security Assertion Markup Language (SAML).
Which solution meets these requirements?

A. Enable AWS 1AM Identity Center (AWS Single Sign-On) between AWS and the on- premises LDAP.
B. Create an 1AM policy that uses AWS credentials, and integrate the policy into LDAP.
C. Set up a process that rotates the I AM credentials whenever LDAP credentials are updated.
D. Develop an on-premises custom identity broker application or process that uses AWS Security Token Service (AWS STS) to get short-lived credentials.

Answer: D

Explanation:
The solution that meets the requirements is to develop an on-premises custom identity broker application or process that uses AWS Security Token Service (AWS STS) to get short-lived credentials. This solution allows the company to use its existing LDAP directory service to authenticate its users to the AWS Management Console, without requiring SAML compatibility. The custom identity broker application or process can act as a proxy between the LDAP directory service and AWS STS, and can request temporary security credentials for the users based on their LDAP attributes and roles. The users can then use these credentials to access the AWS Management Console via a sign-in URL generated by the identity broker. This solution also enhances security by using short-lived credentials that expire after a specified duration.
The other solutions do not meet the requirements because they either require SAML compatibility or do not provide access to the AWS Management Console. Enabling AWS IAM Identity Center (AWS Single Sign-On) between AWS and the on-premises LDAP would require the LDAP directory service to support SAML 2.0, which is not the case for this scenario. Creating an IAM policy that uses AWS credentials and integrating the policy into LDAP would not provide access to the AWS Management Console, but only to the AWS APIs. Setting up a process that rotates the IAM credentials whenever LDAP credentials are updated would also not provide access to the AWS Management Console, but only to the AWS CLI. Therefore, these solutions are not suitable for the given requirements.

NEW QUESTION 9
- (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 10
- (Topic 4)
A company is deploying a new application to Amazon Elastic Kubernetes Service (Amazon EKS) with an AWS Fargate cluster. The application needs a storage solution for data persistence. The solution must be highly available and fault tolerant. The solution also must be shared between multiple application containers. Which solution will meet these requirements with the LEAST operational overhead?

A. Create Amazon Elastic Block Store (Amazon EBS) volumes in the same Availability Zones where EKS worker nodes are place
B. Register the volumes in a StorageClass object on an EKS cluste
C. Use EBS Multi-Attach to share the data between containers.
D. Create an Amazon Elastic File System (Amazon EFS) file syste
E. Register the file system in a StorageClass object on an EKS cluste
F. Use the same file system for all containers.
G. Create an Amazon Elastic Block Store (Amazon EBS) volum
H. Register the volume in a StorageClass object on an EKS cluste
I. Use the same volume for all containers.
J. Create Amazon Elastic File System (Amazon EFS) file systems in the same Availability Zones where EKS worker nodes are place
K. Register the file systems in a StorageClass object on an EKS cluste
L. Create an AWS Lambda function to synchronize the data between file systems.

Answer: B

Explanation:
Amazon EFS is a fully managed, elastic, and scalable file system that can be shared between multiple containers. It provides high availability and fault tolerance by
replicating data across multiple Availability Zones. Amazon EFS is compatible with Amazon EKS and AWS Fargate, and can be registered in a StorageClass object on an EKS cluster. Amazon EBS volumes are not supported by AWS Fargate, and cannot be shared between multiple containers without using EBS Multi-Attach, which has limitations and performance implications. EBS Multi-Attach also requires the volumes to be in the same Availability Zone as the worker nodes, which reduces availability and fault tolerance. Synchronizing data between multiple EFS file systems using AWS Lambda is unnecessary, complex, and prone to errors. References:
? Amazon EFS Storage Classes
? Amazon EKS Storage Classes
? Amazon EBS Multi-Attach

NEW QUESTION 10
- (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 13
- (Topic 4)
A company runs a highly available SFTP service. The SFTP service uses two Amazon EC2
Linux instances that run with elastic IP addresses to accept traffic from trusted IP sources on the internet. The SFTP service is backed by shared storage that is attached to the instances. User accounts are created and managed as Linux users in the SFTP servers.
The company wants a serverless option that provides high IOPS performance and highly configurable security. The company also wants to maintain control over user permissions.
Which solution will meet these requirements?

A. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volum
B. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresse
C. Attach the EBS volume to the SFTP service endpoin
D. Grant users access to the SFTP service.
E. Create an encrypted Amazon Elastic File System (Amazon EFS) volum
F. Create an AWS Transfer Family SFTP service with elastic IP addresses and a VPC endpoint that has internet-facing acces
G. Attach a security group to the endpoint that allows only trusted IP addresse
H. Attach the EFS volume to the SFTP service endpoin
I. Grant users access to the SFTP service.
J. Create an Amazon S3 bucket with default encryption enable
K. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresse
L. Attach the S3 bucket to the SFTP service endpoin
M. Grant users access to the SFTP service.
N. Create an Amazon S3 bucket with default encryption enable
O. Create an AWS Transfer Family SFTP service with a VPC endpoint that has internal access in a private subne
P. Attach a security group that allows only trusted IP addresse
Q. Attach the S3 bucket to the SFTP service endpoin
R. Grant users access to the SFTP service.

Answer: C

Explanation:
AWS Transfer Family is a secure transfer service that enables you to transfer files into and out of AWS storage services using SFTP, FTPS, FTP, and AS2 protocols. You can use AWS Transfer Family to create an SFTP-enabled server with a public endpoint that allows only trusted IP addresses. You can also attach an Amazon S3 bucket with default encryption enabled to the SFTP service endpoint, which will provide high IOPS performance and highly configurable security for your data at rest. You can also maintain control over user permissions by granting users access to the SFTP service using IAM roles or service-managed identities. References: https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html
https://docs.aws.amazon.com/transfer/latest/userguide/create-server-s3.html

NEW QUESTION 15
- (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 17
- (Topic 4)
A company runs an SMB file server in its data center. The file server stores large files that the company frequently accesses for up to 7 days after the file creation date. After 7 days, the company needs to be able to access the files with a maximum retrieval time of 24 hours.
Which solution will meet these requirements?

A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
B. Create an Amazon S3 File Gateway to increase the company's storage spac
C. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
D. Create an Amazon FSx File Gateway to increase the company's storage spac

E. Create an Amazon S3 Lifecycle policy to transition the data after 7 days.
F. Configure access to Amazon S3 for each use
G. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Answer: B

Explanation:
Amazon S3 File Gateway is a service that provides a file-based interface to Amazon S3, which appears as a network file share. It enables you to store and retrieve Amazon S3 objects through standard file storage protocols such as SMB. S3 File Gateway can also cache frequently accessed data locally for low-latency access. S3 Lifecycle policy is a feature that allows you to define rules that automate the management of your objects throughout their lifecycle. You can use S3 Lifecycle policy to transition objects to different storage classes based on their age and access patterns. S3 Glacier Deep Archive is a storage class that offers the lowest cost for long-term data archiving, with a retrieval time of 12 hours or 48 hours. This solution will meet the requirements, as it allows the company to store large files in S3 with SMB file access, and to move the files to S3 Glacier Deep Archive after 7 days for cost savings and compliance.
References:
? 1 provides an overview of Amazon S3 File Gateway and its benefits.
? 2 explains how to use S3 Lifecycle policy to manage object storage lifecycle.
? 3 describes the features and use cases of S3 Glacier Deep Archive storage class.

NEW QUESTION 21
- (Topic 4)
To meet security requirements, a company needs to encrypt all of its application data in transit while communicating with an Amazon RDS MySQL DB instance. A recent security audit revealed that encryption at rest is enabled using AWS Key Management Service (AWS KMS), but data in transit is not enabled.
What should a solutions architect do to satisfy the security requirements?

A. Enable 1AM database authentication on the database.
B. Provide self-signed certificate
C. Use the certificates in all connections to the RDS instance.
D. Take a snapshot of the RDS instanc
E. Restore the snapshot to a new instance with encryption enabled.
F. Download AWS-provided root certificate
G. Provide the certificates in all connections to the RDS instance.

Answer: D

Explanation:
To satisfy the security requirements, the solutions architect should download AWS-provided root certificates and provide the certificates in all connections to the RDS instance. This will enable SSL/TLS encryption for data in transit between the application and the RDS instance. SSL/TLS encryption provides a layer of security by encrypting data that moves between the client and the server. Amazon RDS creates an SSL certificate and installs the certificate on the DB instance when the instance is provisioned. The application can use the AWS-provided root certificates to verify the identity of the DB instance and establish a secure connection1.
The other options are not correct because they do not enable encryption for data in transit or are not relevant for the use case. Enabling IAM database authentication on the database is not correct because this option only provides a method of authentication, not encryption. IAM database authentication allows users to use AWS Identity and Access Management (IAM) users and roles to access a database, instead of using a database user name and password2.
Providing self-signed certificates is not correct because this option is not secure or reliable. Self-signed certificates are certificates that are signed by the same entity that issued them, instead of by a trusted certificate authority (CA). Self-signed certificates can be easily forged or compromised, and are not recognized by most browsers and applications3. Taking a snapshot of the RDS instance and restoring it to a new instance with encryption enabled is not correct because this option only enables encryption at rest, not encryption in transit. Encryption at rest protects data that is stored on disk, but does not protect data that is moving between the client and the server4.
References:
? Using SSL/TLS to encrypt a connection to a DB instance - Amazon Relational Database Service
? IAM database authentication for MySQL and PostgreSQL - Amazon Relational Database Service
? What are self-signed certificates?
? Encrypting Amazon RDS resources - Amazon Relational Database Service

NEW QUESTION 23
- (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 25
- (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 28
- (Topic 4)
A solutions architect is designing a new API using Amazon API Gateway that will receive
requests from users. The volume of requests is highly variable; several hours can pass without receiving a single request. The data processing will take place asynchronously, but should be completed within a few seconds after a request is made.
Which compute service should the solutions architect have the API invoke to deliver the requirements at the lowest cost?

A. An AWS Glue job
B. An AWS Lambda function
C. A containerized service hosted in Amazon Elastic Kubernetes Service (Amazon EKS)
D. A containerized service hosted in Amazon ECS with Amazon EC2

Answer: B

Explanation:
API Gateway + Lambda is the perfect solution for modern applications with serverless architecture.

NEW QUESTION 29
- (Topic 4)
A company needs to create an AWS Lambda function that will run in a VPC in the
company's primary AWS account. The Lambda function needs to access files that the company stores
in an Amazon Elastic File System (Amazon EFS) file system. The EFS file system is located in a secondary AWS account. As the company adds files to the file system the solution must scale to meet the demand.
Which solution will meet these requirements MOST cost-effectively?

A. Create a new EPS file system in the primary account Use AWS DataSync to copy the contents of the original EPS file system to the new EPS file system
B. Create a VPC peering connection between the VPCs that are in the primary account and the secondary account
C. Create a second Lambda function In the secondary account that has a mount that is configured for the file syste
D. Use the primary account's Lambda function to invoke the secondary account's Lambda function
E. Move the contents of the file system to a Lambda Layer's Configure the Lambda layer's permissions to allow the company's secondary account to use the Lambda layer.

Answer: B

Explanation:
This option is the most cost-effective and scalable way to allow the Lambda function in the primary account to access the EFS file system in the secondary account. VPC peering enables private connectivity between two VPCs without requiring gateways, VPN connections, or dedicated network connections. The Lambda function can use the VPC peering connection to mount the EFS file system as a local file system and access the files as needed. The solution does not incur additional data transfer or storage costs, and it leverages the existing EFS file system without duplicating or moving the data.
Option A is not cost-effective because it requires creating a new EFS file system and using AWS DataSync to copy the data from the original EFS file system. This would incur additional storage and data transfer costs, and it would not provide real-time access to the files.
Option C is not scalable because it requires creating a second Lambda function in the secondary account and configuring cross-account permissions to invoke it

from the primary account. This would add complexity and latency to the solution, and it would increase the Lambda invocation costs.
Option D is not feasible because Lambda layers are not designed to store large amounts of data or provide file system access. Lambda layers are used to share common code or libraries across multiple Lambda functions. Moving the contents of the EFS file system to a Lambda layer would exceed the size limit of 250 MB for a layer, and it would not allow the Lambda function to read or write files to the layer. References:
? What Is VPC Peering?
? Using Amazon EFS file systems with AWS Lambda
? What Are Lambda Layers?

NEW QUESTION 32
- (Topic 4)
A company operates a two-tier application for image processing. The application uses two Availability Zones, each with one public subnet and one private subnet. An Application Load Balancer (ALB) for the web tier uses the public subnets. Amazon EC2 instances for the application tier use the private subnets.
Users report that the application is running more slowly than expected. A security audit of the web server log files shows that the application is receiving millions of illegitimate requests from a small number of IP addresses. A solutions architect needs to resolve the immediate performance problem while the company investigates a more permanent solution.
What should the solutions architect recommend to meet this requirement?

A. Modify the inbound security group for the web tie
B. Add a deny rule for the IP addresses that are consuming resources.
C. Modify the network ACL for the web tier subnet
D. Add an inbound deny rule for the IP addresses that are consuming resources
E. Modify the inbound security group for the application tie
F. Add a deny rule for the IP addresses that are consuming resources.
G. Modify the network ACL for the application tier subnet
H. Add an inbound deny rule for the IP addresses that are consuming resources

Answer: B

Explanation:
Deny the request from the first entry at the public subnet, dont allow it to cross and get to the private subnet.
In this scenario, the security audit reveals that the application is receiving millions of illegitimate requests from a small number of IP addresses. To address this issue, it is recommended to modify the network ACL (Access Control List) for the web tier subnets. By adding an inbound deny rule specifically targeting the IP addresses that are consuming resources, the network ACL can block the illegitimate traffic at the subnet level before it reaches the web servers. This will help alleviate the excessive load on the web tier and improve the application's performance.

NEW QUESTION 35
- (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 37
- (Topic 4)
A retail company has several businesses. The IT team for each business manages its own AWS account. Each team account is part of an organization in AWS Organizations. Each team monitors its product inventory levels in an Amazon DynamoDB table in the team's own AWS account.
The company is deploying a central inventory reporting application into a shared AWS account. The application must be able to read items from all the teams' DynamoDB tables.
Which authentication option will meet these requirements MOST securely?

A. Integrate DynamoDB with AWS Secrets Manager in the inventory application accoun
B. Configure the application to use the correct secret from Secrets Manager to authenticate and read the DynamoDB tabl
C. Schedule secret rotation for every 30 days.
D. In every business account, create an 1AM user that has programmatic acces
E. Configure the application to use the correct 1AM user access key ID and secret access key to authenticate and read the DynamoDB tabl
F. Manually rotate 1AM access keys every 30 days.
G. In every business account, create an 1AM role named BU_ROLE with a policy that gives the role access to the DynamoDB table and a trust policy to trust a specific role in the inventory application accoun
H. In the inventory account, create a role named APP_ROLE that allows access to the STS AssumeRole API operatio
I. Configure the application to use APP_ROLE and assume the cross-account role BU_ROLE to read the DynamoDB table.
J. Integrate DynamoDB with AWS Certificate Manager (ACM). Generate identity certificates to authenticate DynamoD
K. Configure the application to use the correct certificate to authenticate and read the DynamoDB table.

Answer: C

Explanation:
This solution meets the requirements most securely because it uses IAM roles and the STS AssumeRole API operation to authenticate and authorize the inventory application to access the DynamoDB tables in different accounts. IAM roles are more secure than IAM users or certificates because they do not require long-term credentials or passwords. Instead, IAM roles provide temporary security credentials that are automatically rotated and can be configured with a limited duration. The STS AssumeRole API operation enables you to request temporary credentials for a role that you are allowed to assume. By using this operation, you can delegate access to resources that are in different AWS accounts that you own or that are owned by third parties. The trust policy of the role defines which entities can assume the role, and the permissions policy of the role defines which actions can be performed on the resources. By using this solution, you can avoid hard- coding credentials or certificates in the inventory application, and you can also avoid storing them in Secrets Manager or ACM. You can also leverage the built-in security features of IAM and STS, such as MFA, access logging, and policy conditions.
References:
? IAM Roles
? STS AssumeRole
? Tutorial: Delegate Access Across AWS Accounts Using IAM Roles

NEW QUESTION 42
- (Topic 4)
A solutions architect needs to copy files from an Amazon S3 bucket to an Amazon Elastic File System (Amazon EFS) file system and another S3 bucket. The files must be copied continuously. New files are added to the original S3 bucket consistently. The copied files should be overwritten only if the source file changes.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS DataSync location for both the destination S3 bucket and the EFS file syste
B. Create a task for the destination S3 bucket and the EFS file syste
C. Set the transfer mode to transfer only data that has changed.
D. Create an AWS Lambda functio
E. Mount the file system to the functio
F. Set up an S3 event notification to invoke the function when files are created and changed in Amazon S3. Configure the function to copy files to the file system and the destination S3 bucket.
G. Create an AWS DataSync location for both the destination S3 bucket and the EFS file syste
H. Create a task for the destination S3 bucket and the EFS file syste
I. Set the transfer mode to transfer all data.
J. Launch an Amazon EC2 instance in the same VPC as the file syste
K. Mount the file syste
L. Create a script to routinely synchronize all objects that changed in the origin S3 bucket to the destination S3 bucket and the mounted file system.

Answer: A

Explanation:
AWS DataSync is a service that makes it easy to move large amounts of data between AWS storage services and on-premises storage systems. AWS DataSync can copy files from an S3 bucket to an EFS file system and another S3 bucket continuously, as well as overwrite only the files that have changed in the source. This solution will meet the requirements with the least operational overhead, as it does not require any code development or manual intervention.
References:
? 4 explains how to create AWS DataSync locations for different storage services.
? 5 describes how to create and configure AWS DataSync tasks for data transfer.
? 6 discusses the different transfer modes that AWS DataSync supports.

NEW QUESTION 47
- (Topic 4)
A media company collects and analyzes user activity data on premises. The company wants to migrate this capability to AWS. The user activity data store will continue to grow
and will be petabytes in size. The company needs to build a highly available data ingestion solution that facilitates on-demand analytics of existing data and new data with SQL.
Which solution will meet these requirements with the LEAST operational overhead?

A. Send activity data to an Amazon Kinesis data strea
B. Configure the stream to deliver the data to an Amazon S3 bucket.
C. Send activity data to an Amazon Kinesis Data Firehose delivery strea
D. Configure the stream to deliver the data to an Amazon Redshift cluster.
E. Place activity data in an Amazon S3 bucke
F. Configure Amazon S3 to run an AWS Lambda function on the data as the data arrives in the S3 bucket.
G. Create an ingestion service on Amazon EC2 instances that are spread across multiple Availability Zone
H. Configure the service to forward data to an Amazon RDS Multi-AZ database.

Answer: B

Explanation:
Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This allows you to use your data to gain new insights for your business and customers. The first step to create a data warehouse is to launch a set of nodes, called an Amazon Redshift cluster. After you provision your cluster, you can upload your data set and then perform data analysis queries. Regardless of the size of the data set, Amazon Redshift offers fast query performance using the same SQL-based tools and business intelligence applications that you use today.

NEW QUESTION 52
- (Topic 4)
A company uses Amazon S3 as its data lake. The company has a new partner that must use SFTP to upload data files A solutions architect needs to implement a highly available SFTP solution that minimizes operational overhead.
Which solution will meet these requirements?

A. Use AWS Transfer Family to configure an SFTP-enabled server with a publicly accessible endpoint Choose the S3 data lake as the destination
B. Use Amazon S3 File Gateway as an SFTP server Expose the S3 File Gateway endpoint URL to the new partner Share the S3 File Gateway endpoint with the newpartner

C. Launch an Amazon EC2 instance in a private subnet in a VP
D. Instruct the new partner to upload files to the EC2 instance by using a VP
E. Run a cron job script on the EC2 instance to upload files to the S3 data lake
F. Launch Amazon EC2 instances in a private subnet in a VP
G. Place a Network Load Balancer (NLB) in front of the EC2 instance
H. Create an SFTP listener port for the NLBShare the NLB hostname with the new partner Run a cron job script on the EC2 instances to upload files to the S3 data lake.

Answer: A

Explanation:
This option is the most cost-effective and simple way to enable SFTP access to the S3 data lake. AWS Transfer Family is a fully managed service that supports secure file transfers over SFTP, FTPS, and FTP protocols. You can create an SFTP-enabled server with a public endpoint and associate it with your S3 bucket. You can also use AWS Identity and Access Management (IAM) roles and policies to control access to your S3 data lake. The service scales automatically to handle any volume of file transfers and provides high availability and durability. You do not need to provision, manage, or patch any servers or load balancers.
Option B is not correct because Amazon S3 File Gateway is not an SFTP server. It is a hybrid cloud storage service that provides a local file system interface to S3. You can use it to store and retrieve files as objects in S3 using standard file protocols such as NFS and SMB. However, it does not support SFTP protocol, and it requires deploying a file gateway appliance on-premises or on EC2.
Option C is not cost-effective or scalable because it requires launching and managing an EC2 instance in a private subnet and setting up a VPN connection for the new partner. This would incur additional costs for the EC2 instance, the VPN connection, and the data transfer. It would also introduce complexity and security risks to the solution. Moreover, it would require running a cron job script on the EC2 instance to upload files to the S3 data lake, which is not efficient or reliable.
Option D is not cost-effective or scalable because it requires launching and managing multiple EC2 instances in a private subnet and placing a NLB in front of them. This would incur additional costs for the EC2 instances, the NLB, and the data transfer. It would also introduce complexity and security risks to the solution. Moreover, it would require running a cron job script on the EC2 instances to upload files to the S3 data lake, which is not efficient or reliable. References:
? What Is AWS Transfer Family?
? What Is Amazon S3 File Gateway?
? What Is Amazon EC2?
? [What Is Amazon Virtual Private Cloud?]
? [What Is a Network Load Balancer?]

NEW QUESTION 54
- (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 57
- (Topic 4)
A company's reporting system delivers hundreds of .csv files to an Amazon S3 bucket each day. The company must convert these files to Apache Parquet format and must store the files in a transformed data bucket.
Which solution will meet these requirements with the LEAST development effort?

A. Create an Amazon EMR cluster with Apache Spark installe
B. Write a Spark application to transform the dat
C. Use EMR File System (EMRFS) to write files to the transformed data bucket.
D. Create an AWS Glue crawler to discover the dat
E. Create an AWS Glue extract, transform, and load (ETL) job to transform the dat
F. Specify the transformed data bucket in the output step.
G. Use AWS Batch to create a job definition with Bash syntax to transform the data and output the data to the transformed data bucke
H. Use the job definition to submit a jo
I. Specify an array job as the job type.
J. Create an AWS Lambda function to transform the data and output the data to the transformed data bucke
K. Configure an event notification for the S3 bucke
L. Specify the Lambda function as the destination for the event notification.

Answer: B

Explanation:
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/three- aws-glue-etl-job-types-for-converting-data-to-apache-parquet.html

NEW QUESTION 60
- (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 64
- (Topic 4)
An ecommerce company stores terabytes of customer data in the AWS Cloud. The data contains personally identifiable information (Pll). The company wants to use the data in three applications. Only one of the applications needs to process the Pll. The Pll must be removed before the other two applications process the data.
Which solution will meet these requirements with the LEAST operational overhead?

A. Store the data in an Amazon DynamoDB tabl
B. Create a proxy application layer to intercept and process the data that each application requests.
C. Store the data in an Amazon S3 bucke
D. Process and transform the data by using S3 Object Lambda before returning the data to the requesting application.
E. Process the data and store the transformed data in three separate Amazon S3 buckets so that each application has its own custom datase
F. Point each application to its respectiveS3 bucket.
G. Process the data and store the transformed data in three separate Amazon DynamoDB tables so that each application has its own custom datase
H. Point each application to its respective DynamoDB table.

Answer: B

Explanation:
https://aws.amazon.com/blogs/aws/introducing-amazon-s3-object-lambda- use-your-code-to-process-data-as-it-is-being-retrieved-from-s3/
S3 Object Lambda is a new feature of Amazon S3 that enables customers to add their own code to process data retrieved from S3 before returning it to the application. By using S3 Object Lambda, the data can be processed and transformed in real-time, without the need to store multiple copies of the data in separate S3 buckets or DynamoDB tables.
In this case, the Pll can be removed from the data by the code added to S3 Object Lambda before returning the data to the two applications that do not need to process Pll. The one application that requires Pll can be pointed to the original S3 bucket where the Pll is still stored.
Using S3 Object Lambda is the simplest and most cost-effective solution, as it eliminates the need to maintain multiple copies of the same data in different buckets or tables, which can result in additional storage costs and operational overhead.

NEW QUESTION 66
- (Topic 4)
A company is running a microservices application on Amazon EC2 instances. The company wants to migrate the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for scalability. The company must configure the Amazon EKS control plane with endpoint private access set to true and endpoint public access set to false to maintain security compliance The company must also put the data plane in private subnets. However, the company has received error notifications because the node cannot join the cluster.
Which solution will allow the node to join the cluster?

A. Grant the required permission in AWS Identity and Access Management (1AM) to the AmazonEKSNodeRole 1AM role.
B. Create interface VPC endpoints to allow nodes to access the control plane.
C. Recreate nodes in the public subnet Restrict security groups for EC2 nodes
D. Allow outbound traffic in the security group of the nodes.

Answer: B

Explanation:
Kubernetes API requests within your cluster's VPC (such as node to control plane communication) use the private VPC endpoint.
https://docs.aws.amazon.com/eks/latest/userguide/cluster-endpoint.html

NEW QUESTION 71
- (Topic 4)
The customers of a finance company request appointments with financial advisors by sending text messages. A web application that runs on Amazon EC2 instances accepts the appointment requests. The text messages are published to an Amazon Simple Queue Service (Amazon SQS) queue through the web application. Another application that runs on EC2 instances then sends meeting invitations and meeting confirmation email messages to the customers. After

successful scheduling, this application stores the meeting information in an Amazon DynamoDB database.
As the company expands, customers report that their meeting invitations are taking longer to arrive.
What should a solutions architect recommend to resolve this issue?

A. Add a DynamoDB Accelerator (DAX) cluster in front of the DynamoDB database.
B. Add an Amazon API Gateway API in front of the web application that accepts the appointment requests.
C. Add an Amazon CloudFront distributio
D. Set the origin as the web application that accepts the appointment requests.
E. Add an Auto Scaling group for the application that sends meeting invitation
F. Configure the Auto Scaling group to scale based on the depth of the SQS queue.

Answer: D

Explanation:
To resolve the issue of longer delivery times for meeting invitations, the solutions architect can recommend adding an Auto Scaling group for the application that sends meeting invitations and configuring the Auto Scaling group to scale based on the depth of the SQS queue. This will allow the application to scale up as the number of appointment requests increases, improving the performance and delivery times of the meeting invitations.

NEW QUESTION 74
- (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 75
- (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 76
- (Topic 4)
A company runs multiple Amazon EC2 Linux instances in a VPC across two Availability Zones. The instances host applications that use a hierarchical directory structure. The applications need to read and write rapidly and concurrently to shared storage.
What should a solutions architect do to meet these requirements?

A. Create an Amazon S3 bucke
B. Allow access from all the EC2 instances in the VPC.
C. Create an Amazon Elastic File System (Amazon EFS) file syste
D. Mount the EFS file system from each EC2 instance.
E. Create a file system on a Provisioned IOPS SSD (102) Amazon Elastic Block Store (Amazon EBS) volum

F. Attach the EBS volume to all the EC2 instances.
G. Create file systems on Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instanc
H. Synchromze the EBS volumes across the different EC2 instances.

Answer: B

Explanation:
it allows the EC2 instances to read and write rapidly and concurrently to shared storage across two Availability Zones. Amazon EFS provides a scalable, elastic, and highly available file system that can be mounted from multiple EC2 instances. Amazon EFS supports high levels of throughput and IOPS, and consistent low latencies. Amazon EFS also supports NFSv4 lock upgrading and downgrading, which enables high levels of concurrency. References:
? Amazon EFS Features
? Using Amazon EFS with Amazon EC2

NEW QUESTION 78
- (Topic 4)
A company is deploying an application in three AWS Regions using an Application Load Balancer Amazon Route 53 will be used to distribute traffic between these Regions. Which Route 53 configuration should a solutions architect use to provide the MOST high- performing experience?

A. Create an A record with a latency policy.
B. Create an A record with a geolocation policy.
C. Create a CNAME record with a failover policy.
D. Create a CNAME record with a geoproximity policy.

Answer: A

Explanation:
To provide the most high-performing experience for the users of the application, a solutions architect should use a latency routing policy for the Route 53 A record. This policy allows Route 53 to route traffic to the AWS Region that provides the lowest possible latency for the users1. A latency routing policy can also improve the availability of the application, as Route 53 can automatically route traffic to another Region if the primary Region becomes unavailable2.
References:
? 1: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing- policy.html#routing-policy-latency
? 2: https://aws.amazon.com/route53/faqs/#Latency_Based_Routing

NEW QUESTION 79
- (Topic 4)
A company has two VPCs that are located in the us-west-2 Region within the same AWS account. The company needs to allow network traffic between these VPCs. Approximately 500 GB of data transfer will occur between the VPCs each month.
What is the MOST cost-effective solution to connect these VPCs?

A. Implement AWS Transit Gateway to connect the VPC
B. Update the route tables of each VPC to use the transit gateway for inter-VPC communication.
C. Implement an AWS Site-to-Site VPN tunnel between the VPC
D. Update the route tables of each VPC to use the VPN tunnel for inter-VPC communication.
E. Set up a VPC peering connection between the VPC
F. Update the route tables of each VPC to use the VPC peering connection for inter-VPC communication.
G. Set up a 1 GB AWS Direct Connect connection between the VPC
H. Update the route tables of each VPC to use the Direct Connect connection for inter-VPC communication.

Answer: C

Explanation:
To connect two VPCs in the same Region within the same AWS account, VPC peering is the most cost-effective solution. VPC peering allows direct network traffic between the VPCs without requiring a gateway, VPN connection, or AWS Transit Gateway. VPC peering also does not incur any additional charges for data transfer between the VPCs.
References:
? What Is VPC Peering?
? VPC Peering Pricing

NEW QUESTION 80
- (Topic 4)
A solutions architect must provide an automated solution for a company's compliance policy that states security groups cannot include a rule that allows SSH from 0.0.0.0/0. The company needs to be notified if there is any breach in the policy. A solution is needed as soon as possible.
What should the solutions architect do to meet these requirements with the LEAST operational overhead?

A. Write an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one.
B. Enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created.
C. Create an 1AM role with permissions to globally open security groups and network ACL
D. Create an Amazon Simple Notification Service (Amazon SNS) topic to generate a notification every time the role is assumed by a user.
E. Configure a service control policy (SCP) that prevents non-administrative users from creating or editing security group
F. Create a notification in the ticketing system when a user requests a rule that needs administrator permissions.

Answer: B

Explanation:
The most suitable solution for the company’s compliance policy is to enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created. This solution has the least operational overhead because it uses a predefined rule that is already available in AWS Config, which is a service that enables users to assess, audit, and evaluate the configurations of their AWS resources. The restricted-ssh rule checks whether security groups that are in use have inbound rules that allow SSH from 0.0.0.0/0 addresses, and reports them as noncompliant1. Users can configure the rule to send notifications to an Amazon SNS topic when a noncompliant change occurs, and subscribe to the topic to

receive alerts via email, SMS, or other methods2.
The other options are not correct because they either have more operational overhead or do not meet the requirements. Writing an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one is not correct because it requires custom code development and maintenance, which adds complexity and cost to the solution. Creating an IAM role with permissions to globally open security groups and network ACLs, and creating an Amazon SNS topic to generate a notification every time the role is assumed by a user is not correct because it does not prevent or detect the creation of noncompliant rules by other users or roles, and it does not address the existing rules that may violate the policy. Configuring a service control policy (SCP) that prevents non-administrative users from creating or editing security groups, and creating a notification in the ticketing system when a user requests a rule that needs administrator permissions is not correct because it does not provide an automated solution for the policy enforcement and notification, and it may limit the flexibility and productivity of the users.
References:
? restricted-ssh - AWS Config
? Getting Notifications When Your Resources Change - AWS Config

NEW QUESTION 82
- (Topic 4)
A company runs a real-time data ingestion solution on AWS. The solution consists of the most recent version of Amazon Managed Streaming for Apache Kafka (Amazon MSK). The solution is deployed in a VPC in private subnets across three Availability Zones.
A solutions architect needs to redesign the data ingestion solution to be publicly available over the internet. The data in transit must also be encrypted.
Which solution will meet these requirements with the MOST operational efficiency?

A. Configure public subnets in the existing VP
B. Deploy an MSK cluster in the public subnet
C. Update the MSK cluster security settings to enable mutual TLS authentication.
D. Create a new VPC that has public subnet
E. Deploy an MSK cluster in the public subnet
F. Update the MSK cluster security settings to enable mutual TLS authentication.
G. Deploy an Application Load Balancer (ALB) that uses private subnet
H. Configure an ALB security group inbound rule to allow inbound traffic from the VPC CIDR block for HTTPS protocol.
I. Deploy a Network Load Balancer (NLB) that uses private subnet
J. Configure an NLB listener for HTTPS communication over the internet.

Answer: A

Explanation:
The solution that meets the requirements with the most operational efficiency is to configure public subnets in the existing VPC and deploy an MSK cluster in the public subnets. This solution allows the data ingestion solution to be publicly available over the internet without creating a new VPC or deploying a load balancer. The solution also ensures that the data in transit is encrypted by enabling mutual TLS authentication, which requires both the client and the server to present certificates for verification. This solution leverages the public access feature of Amazon MSK, which is available for clusters running Apache Kafka 2.6.0 or later versions1.
The other solutions are not as efficient as the first one because they either create unnecessary resources or do not encrypt the data in transit. Creating a new VPC with public subnets would incur additional costs and complexity for managing network resources and routing. Deploying an ALB or an NLB would also add more costs and latency for the data ingestion solution. Moreover, an ALB or an NLB would not encrypt the data in transit by itself, unless they are configured with HTTPS listeners and certificates, which would require additional steps and maintenance. Therefore, these solutions are not optimal for the given requirements.
References:
? Public access - Amazon Managed Streaming for Apache Kafka

NEW QUESTION 84
- (Topic 4)
A company needs a solution to prevent photos with unwanted content from being uploaded to the company's web application. The solution must not involve training a machine learning (ML) model. Which solution will meet these requirements?

A. Create and deploy a model by using Amazon SageMaker Autopilo
B. Create a real-time endpoint that the web application invokes when new photos are uploaded.
C. Create an AWS Lambda function that uses Amazon Rekognition to detect unwanted conten
D. Create a Lambda function URL that the web application invokes when new photos are uploaded.
E. Create an Amazon CloudFront function that uses Amazon Comprehend to detect unwanted conten
F. Associate the function with the web application.
G. Create an AWS Lambda function that uses Amazon Rekognition Video to detectunwanted conten
H. Create a Lambda function URL that the web application invokes when new photos are uploaded.

Answer: B

Explanation:
The solution that will meet the requirements is to create an AWS Lambda function that uses Amazon Rekognition to detect unwanted content, and create a Lambda function URL that the web application invokes when new photos are uploaded. This solution does not involve training a machine learning model, as Amazon Rekognition is a fully managed service that provides pre-trained computer vision models for image and video analysis. Amazon Rekognition can detect unwanted content such as explicit or suggestive adult content, violence, weapons, drugs, and more. By using AWS Lambda, the company can create a serverless function that can be triggered by an HTTP request from the web application. The Lambda function can use the Amazon Rekognition API to analyze the uploaded photos and return a response indicating whether they contain unwanted content or not.
The other solutions are not as effective as the first one because they either involve training a machine learning model, do not support image analysis, or do not work with photos. Creating and deploying a model by using Amazon SageMaker Autopilot involves training a machine learning model, which is not required for the scenario. Amazon SageMaker Autopilot is a service that automatically creates, trains, and tunes the best machine learning models for classification or regression based on the data provided by the user. Creating an Amazon CloudFront function that uses Amazon Comprehend to detect unwanted content does not support image analysis, as Amazon Comprehend is a natural language processing service that analyzes text, not images. Amazon Comprehend can extract insights and relationships from text such as language, sentiment, entities, topics, and more. Creating an AWS Lambda function that uses Amazon Rekognition Video to detect unwanted content does not work with photos, as Amazon Rekognition Video is designed for analyzing video streams, not static images. Amazon Rekognition Video can detect activities, objects, faces, celebrities, text, and more in video streams.
References:
? Amazon Rekognition
? AWS Lambda
? Detecting unsafe content - Amazon Rekognition

? Amazon SageMaker Autopilot
? Amazon Comprehend

NEW QUESTION 86
- (Topic 4)
A company manages AWS accounts in AWS Organizations. AWS 1AM Identity Center (AWS Single Sign-On) and AWS Control Tower are configured for the accounts. The company wants to manage multiple user permissions across all the accounts.
The permissions will be used by multiple 1AM users and must be split between the developer and administrator teams. Each team requires different permissions.
The company wants a solution that includes new users that are hired on both teams.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create individual users in 1AM Identity Center (or each accoun
B. Create separate developer and administrator groups in 1AM Identity Cente
C. Assign the users to the appropriate groups Create a custom 1AM policy for each group to set fine-grained permissions.
D. Create individual users in 1AM Identity Center for each accoun
E. Create separate developer and administrator groups in 1AM Identity Cente
F. Assign the users to the appropriate group
G. Attach AWS managed 1AM policies to each user as needed for fine- grained permissions.
H. Create individual users in 1AM Identity Center Create new developer and administrator groups in 1AM Identity Cente
I. Create new permission sets that include the appropriate 1AM policies for each grou
J. Assign the new groups to the appropriate accounts Assign the new permission sets to the new groups When new users are hired, add them to the appropriate group.
K. Create individual users in 1AM Identity Cente
L. Create new permission sets that include the appropriate 1AM policies for each use
M. Assign the users to the appropriate account
N. Grant additional 1AM permissions to the users from within specific account
O. When new users are hired, add them to 1AM Identity Center and assign them to the accounts.

Answer: C

Explanation:
This solution meets the requirements with the least operational overhead because it leverages the features of IAM Identity Center and AWS Control Tower to centrally manage multiple user permissions across all the accounts. By creating new groups and permission sets, the company can assign fine-grained permissions to the developer and administrator teams based on their roles and responsibilities. The permission sets are applied to the groups at the organization level, so they are automatically inherited by all the accounts in the organization. When new users are hired, the company only needs to add them to the appropriate group in IAM Identity Center, and they will automatically get the permissions assigned to that group. This simplifies the user management and reduces the manual effort of assigning permissions to each user individually.
References:
? Managing access to AWS accounts and applications
? Managing permissions sets
? Managing groups

NEW QUESTION 91
- (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 95
- (Topic 4)
A company runs container applications by using Amazon Elastic Kubernetes Service (Amazon EKS). The company's workload is not consistent throughout the day
The company wants Amazon EKS to scale in and out according to the workload.

Which combination of steps will meet these requirements with the LEAST operational overhead? {Select TWO.)

A. Use an AWS Lambda function to resize the EKS cluster
B. Use the Kubernetes Metrics Server to activate horizontal pod autoscaling.
C. Use the Kubernetes Cluster Autoscaler to manage the number of nodes in the cluster.
D. Use Amazon API Gateway and connect it to Amazon EKS
E. Use AWS App Mesh to observe network activity.

Answer: BC

Explanation:
https://docs.aws.amazon.com/eks/latest/userguide/horizontal-pod- autoscaler.html https://docs.aws.amazon.com/eks/latest/userguide/autoscaling.html
Horizontal pod autoscaling is a feature of Kubernetes that automatically scales the number of pods in a deployment, replication controller, or replica set based on that resource's CPU utilization. It requires a metrics source such as the Kubernetes Metrics Server to provide CPU usage data1. Cluster autoscaling is a feature of Kubernetes that automatically adjusts the number of nodes in a cluster when pods fail or are rescheduled onto other nodes. It requires an integration with AWS Auto Scaling groups to manage the EC2 instances that join the cluster2. By using both horizontal pod autoscaling and cluster autoscaling, the solution can ensure that Amazon EKS scales in and out according to the workload.

NEW QUESTION 96
- (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 97
- (Topic 4)
A company runs a three-tier web application in the AWS Cloud that operates across three Availability Zones. The application architecture has an Application Load Balancer, an Amazon EC2 web server that hosts user session states, and a MySQL database that runs on an EC2 instance. The company expects sudden increases in application traffic. The company wants to be able to scale to meet future application capacity demands and to ensure high availability across all three Availability Zones.
Which solution will meet these requirements?

A. Migrate the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deploymen
B. Use Amazon ElastiCache for Redis with high availability to store session data and to cache read
C. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
D. Migrate the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deploymen
E. Use Amazon ElastiCache for Memcached with high availability to store session data and to cache read
F. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
G. Migrate the MySQL database to Amazon DynamoD
H. Use DynamoDB Accelerator (DAX) to cache read
I. Store the session data in DynamoD
J. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
K. Migrate the MySQL database to Amazon RDS for MySQL in a single Availability Zon
L. Use Amazon ElastiCache for Redis with high availability to store session data and to cache read
M. Migrate the web server to an Auto Scaling group that is in three Availability Zones.

Answer: A

Explanation:
This answer is correct because it meets the requirements of scaling to meet future application capacity demands and ensuring high availability across all three Availability Zones. By migrating the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deployment, the company can benefit from automatic

failover, backup, and patching of the database across multiple Availability Zones. By using Amazon ElastiCache for Redis with high availability, the company can store session data and cache reads in a fast, in-memory data store that can also fail over across Availability Zones. By migrating the web server to an Auto Scaling group that is in three Availability Zones, the company can automatically scale the web server capacity based on the demand and traffic
patterns. References:
? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.ht
ml
? https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html
? https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2- auto-scaling.html

NEW QUESTION 98
- (Topic 4)
A company uses Amazon Elastic Kubernetes Service (Amazon EKS) to run a container application. The EKS cluster stores sensitive information in the Kubernetes secrets object. The company wants to ensure that the information is encrypted
Which solution will meet these requirements with the LEAST operational overhead?

A. Use the container application to encrypt the information by using AWS Key Management Service (AWS KMS).
B. Enable secrets encryption in the EKS cluster by using AWS Key Management Service (AWS KMS)_
C. Implement an AWS Lambda tuncüon to encrypt the information by using AWS Key Management Service (AWS KMS).
D. use AWS Systems Manager Parameter Store to encrypt the information by using AWS Key Management Service (AWS KMS).

Answer: B

Explanation:
it allows the company to encrypt the Kubernetes secrets object in the EKS cluster with the least operational overhead. By enabling secrets encryption in the EKS cluster, the company can use AWS Key Management Service (AWS KMS) to generate and manage encryption keys for encrypting and decrypting secrets at rest. This is a simple and secure way to protect sensitive information in EKS clusters. References:
? Encrypting Kubernetes secrets with AWS KMS
? Kubernetes Secrets

NEW QUESTION 101
- (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 102
- (Topic 4)
A company's applications run on Amazon EC2 instances in Auto Scaling groups. The company notices that its applications experience sudden traffic increases on random days of the week The company wants to maintain application performance during sudden traffic increases.
Which solution will meet these requirements MOST cost-effectively?

A. Use manual scaling to change the size of the Auto Scaling group.
B. Use predictive scaling to change the size of the Auto Scaling group.
C. Use dynamic scaling to change the size of the Auto Scaling group.
D. Use schedule scaling to change the size of the Auto Scaling group

Answer: C

Explanation:
Dynamic scaling is a type of autoscaling that automatically adjusts the number of EC2 instances in an Auto Scaling group based on demand or load. It uses CloudWatch alarms to trigger scaling actions when a specified metric crosses a threshold. It can scale out (add instances) or scale in (remove instances) as neede1d. By using dynamic scaling, the solution can maintain application performance during sudden traffic increases most cost-effectively.
* A. Use manual scaling to change the size of the Auto Scaling group. This solution will not meet the requirement of maintaining application performance during sudden traffic increases, as manual scaling requires users to manually increase or decrease the number of instances through a CLI or console. It does not respond automatically to changes in demand or load2.
* B. Use predictive scaling to change the size of the Auto Scaling group. This solution will not meet the requirement of most cost-effectiveness, as predictive scaling uses machine learning and artificial intelligence tools to evaluate traffic loads and anticipate when more or fewer resources are needed. It performs scheduled scaling actions based on the prediction, which may not match the actual demand or load at any given time. Predictive scaling is more suitable for scenarios where there are predictable traffic patterns or known changes in traffic loads3.
* D. Use schedule scaling to change the size of the Auto Scaling group. This solution will not meet the requirement of maintaining application performance during sudden traffic increases, as schedule scaling performs scaling actions at specific times that users schedule. It does not respond automatically to changes in demand or load. Schedule scaling is more suitable for scenarios where there are predictable traffic drops or spikes at specific times of the day.
Reference URL: https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scale-based-on-demand.html

NEW QUESTION 104
- (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 109
- (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 112
- (Topic 4)
A company is preparing a new data platform that will ingest real-time streaming data from multiple sources. The company needs to transform the data before writing the data to Amazon S3. The company needs the ability to use SQL to query the transformed data.
Which solutions will meet these requirements? (Choose two.)

A. Use Amazon Kinesis Data Streams to stream the dat
B. Use Amazon Kinesis Data Analytics to transform the dat
C. Use Amazon Kinesis Data Firehose to write the data to Amazon S3. Use Amazon Athena to query the transformed data from Amazon S3.
D. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to stream the dat
E. Use AWS Glue to transform the data and to write the data to Amazon S3. Use Amazon Athena to query the transformed data from Amazon S3.
F. Use AWS Database Migration Service (AWS DMS) to ingest the dat
G. Use Amazon EMR to transform the data and to write the data to Amazon S3. Use Amazon Athena to query the transformed data from Amazon S3.
H. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to stream the dat
I. Use Amazon Kinesis Data Analytics to transform the data and to write the data to Amazon S3. Use the Amazon RDS query editor to query the transformed data from Amazon S3.
J. Use Amazon Kinesis Data Streams to stream the dat
K. Use AWS Glue to transform the dat
L. Use Amazon Kinesis Data Firehose to write the data to Amazon S3. Use the Amazon RDS query editor to query the transformed data from Amazon S3.

Answer: AB

Explanation:
To ingest, transform, and query real-time streaming data from multiple sources, Amazon Kinesis and Amazon MSK are suitable solutions. Amazon Kinesis Data Streams can stream the data from various sources and integrate with other AWS services. Amazon Kinesis Data Analytics can transform the data using SQL or Apache Flink.
Amazon Kinesis Data Firehose can write the data to Amazon S3 or other destinations. Amazon Athena can query the transformed data from Amazon S3 using standard SQL. Amazon MSK can stream the data using Apache Kafka, which is a popular open-source platform for streaming data. AWS Glue can transform the data using Apache Spark or Python scripts and write the data to Amazon S3 or other destinations. Amazon Athena can also query the transformed data from Amazon S3 using standard SQL.
References:
? What Is Amazon Kinesis Data Streams?
? What Is Amazon Kinesis Data Analytics?
? What Is Amazon Kinesis Data Firehose?
? What Is Amazon Athena?
? What Is Amazon MSK?
? What Is AWS Glue?

NEW QUESTION 116
- (Topic 4)

A company has an organization in AWS Organizations. The company runs Amazon EC2 instances across four AWS accounts in the root organizational unit (OU). There are three nonproduction accounts and one production account. The company wants to prohibit users from launching EC2 instances of a certain size in the nonproduction accounts. The company has created a service control policy (SCP) to deny access to launch instances that use the prohibited types.
Which solutions to deploy the SCP will meet these requirements? (Select TWO.)

A. Attach the SCP to the root OU for the organization.
B. Attach the SCP to the three nonproduction Organizations member accounts.
C. Attach the SCP to the Organizations management account.
D. Create an OU for the production accoun
E. Attach the SCP to the O
F. Move the production member account into the new OU.
G. Create an OU for the required account
H. Attach the SCP to the O
I. Move the nonproduction member accounts into the new OU.

Answer: BE

Explanation:
SCPs are a type of organization policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization. SCPs help you to ensure your accounts stay within your organization's access control guidelines1.
To apply an SCP to a specific set of accounts, you need to create an OU for those accounts and attach the SCP to the OU. This way, the SCP affects only the member accounts in that OU and not the other accounts in the organization. If you attach the SCP to the root OU, it will apply to all accounts in the organization, including the production account, which is not the desired outcome. If you attach the SCP to the management account, it will have no effect, as SCPs do not affect users or roles in the management
account1.
Therefore, the best solutions to deploy the SCP are B and E. Option B attaches the SCP directly to the three nonproduction accounts, while option E creates a separate OU for the nonproduction accounts and attaches the SCP to the OU. Both options will achieve the same result of restricting the EC2 instance types in the nonproduction accounts, but option E might be more scalable and manageable if there are more accounts or policies to be applied in the future2.
References:
? 1: Service control policies (SCPs) - AWS Organizations
? 2: Best Practices for AWS Organizations Service Control Policies in a Multi- Account Environment

NEW QUESTION 121
- (Topic 4)
A company has multiple AWS accounts for development work. Some staff consistently use oversized Amazon EC2 instances, which causes the company to exceed the yearly budget for the development accounts The company wants to centrally restrict the creation of AWS resources in these accounts
Which solution will meet these requirements with the LEAST development effort?

A. Develop AWS Systems Manager templates that use an approved EC2 creation proces
B. Use the approved Systems Manager templates to provision EC2 instances.
C. Use AWS Organizations to organize the accounts into organizational units (OUs). Define and attach a service control policy (SCP) to control the usage of EC2 instance types.
D. Configure an Amazon EventBridge rule that invokes an AWS Lambda function when an EC2 instance is create
E. Stop disallowed EC2 instance types.
F. Set up AWS Service Catalog products for the staff to create the allowed EC2 instance types Ensure that staff can deploy EC2 instances only by using the Service Catalog products.

Answer: B

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 using AWS Organizations, the solution can centrally restrict the creation of AWS resources in the development accounts.
* A. Develop AWS Systems Manager templates that use an approved EC2 creation process. Use the approved Systems Manager templates to provision EC2 instances. This solution will not meet the requirement of the least development effort, as it involves developing and maintaining custom templates for EC2 creation, and relying on the staff to use the approved templates instead of enforcing a restriction2.
* C. Configure an Amazon EventBridge rule that invokes an AWS Lambda function when an EC2 instance is created. Stop disallowed EC2 instance types. This solution will not meet the requirement of the least development effort, as it involves writing custom code for Lambda functions, and handling events and errors for EC2 creation3.
* D. Set up AWS Service Catalog products for the staff to create the allowed EC2 instance types En-sure that staff can deploy EC2 instances only by using the Service Catalog products. This solution will not meet the requirement of the least development effort, as it involves setting up and managing Service Catalog products for EC2 creation, and ensuring that staff can only use Service Catalog products instead of enforcing a restriction. Reference URL: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps. html

NEW QUESTION 123
- (Topic 4)
A company needs to extract the names of ingredients from recipe records that are stored as text files in an Amazon S3 bucket A web application will use the ingredient names to query an Amazon DynamoDB table and determine a nutrition score.
The application can handle non-food records and errors The company does not have any employees who have machine learning knowledge to develop this solution
Which solution will meet these requirements MOST cost-effectively?

A. Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur Program the Lambda function to analyze the object and extract the ingredient names by using Amazon Comprehend Store the Amazon Comprehend output in the DynamoDB table.
B. Use an Amazon EventBridge rule to invoke an AWS Lambda function when PutObject requests occu
C. Program the Lambda function to analyze the object by using Amazon Forecast to extract the ingredient names Store the Forecast output in the DynamoDB table.
D. Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur Use Amazon Polly to create audio recordings of the recipe record
E. Save the audio files in the S3 bucket Use Amazon Simple Notification Service (Amazon SNS) to send a URL as a message to employees Instruct the

employees to listen to the audio files and calculate the nutrition score Store the ingredient names in the DynamoDB table.
F. Use an Amazon EventBridge rule to invoke an AWS Lambda function when a PutObject request occurs Program the Lambda function to analyze the object and extract the ingredient names by using Amazon SageMaker Store the inference output from the SageMaker endpoint in the DynamoDB table.

Answer: A

Explanation:
This solution meets the following requirements:
? It is cost-effective, as it only uses serverless components that are charged based on usage and do not require any upfront provisioning or maintenance.
? It is scalable, as it can handle any number of recipe records that are uploaded to the S3 bucket without any performance degradation or manual intervention.
? It is easy to implement, as it does not require any machine learning knowledge or complex data processing logic. Amazon Comprehend is a natural language processing service that can automatically extract entities such as ingredients from text files. The Lambda function can simply invoke the Comprehend API and store the results in the DynamoDB table.
? It is reliable, as it can handle non-food records and errors gracefully. Amazon Comprehend can detect the language and domain of the text files and return an appropriate response. The Lambda function can also implement error handling and logging mechanisms to ensure the data quality and integrity.
References:
? Using AWS Lambda with Amazon S3 - AWS Lambda
? What Is Amazon Comprehend? - Amazon Comprehend
? Working with Tables - Amazon DynamoDB

NEW QUESTION 126
- (Topic 4)
A group requires permissions to list an Amazon S3 bucket and delete objects from that bucket An administrator has created the following 1AM policy to provide access to the bucket and applied that policy to the group. The group is not able to delete objects in the bucket. The company follows least-privilege access rules.
A)

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "s3:ListBucket",
                "s3:DeleteObject"
            ],
            "Resource": [
                "arn:aws:s3:::bucket-name"
            ],
            "Effect": "Allow"
        }
    ]
}
```

B)

```
"Action": [
    "s3:*Object"
],
"Resource": [
    "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

C)

```
"Action": [
    "s3:DeleteObject"
],
"Resource": [
    "arn:aws:s3:::bucket-name*"
],
"Effect": "Allow"
```

D)

```
"Action": [
    "s3:DeleteObject"
],
"Resource": [
    "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

A. Option A
B. Option B
C. Option C

D. Option D

Answer: D

Explanation:
{ "Version": "2012-10-17",
"Statement": [
{
"Action": ["s3:ListBucket", "s3:DeleteObject"
],
"Resource": ["arn:aws:s3:::<bucket-name>"
],
"Effect": "Allow",
},
{
"Action": "s3:*DeleteObject", "Resource": [
"arn:aws:s3:::<bucket-name>/*" # <- The policy clause kludge "added" to match the solution (Q248.1) example
],
"Effect": "Allow"
}
]
}

NEW QUESTION 131
- (Topic 4)
A company has 150 TB of archived image data stored on-premises that needs to be moved to the AWS Cloud within the next month. The company's current network connection allows up to 100 Mbps uploads for this purpose during the night only.
What is the MOST cost-effective mechanism to move this data and meet the migration deadline?

A. Use AWS Snowmobile to ship the data to AWS.
B. Order multiple AWS Snowball devices to ship the data to AWS.
C. Enable Amazon S3 Transfer Acceleration and securely upload the data.
D. Create an Amazon S3 VPC endpoint and establish a VPN to upload the data.

Answer: B

Explanation:
AWS Snowball is a petabyte-scale data transport service that uses secure devices to transfer large amounts of data into and out of the AWS Cloud. Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns. AWS Snowball can transfer up to 80 TB of data per device, and multiple devices can be used in parallel to meet the migration deadline. AWS Snowball is more cost-effective than AWS Snowmobile, which is designed for exabyte-scale data transfers, or Amazon S3 Transfer Acceleration, which is optimized for fast transfers over long distances. Amazon S3 VPC endpoint does not increase the upload speed, but only provides a secure and private connection between the VPC and
S3. References: AWS Snowball, AWS Snowmobile, Amazon S3 Transfer Acceleration, Amazon S3 VPC endpoint

NEW QUESTION 134
- (Topic 4)
A solutions architect needs to review a company's Amazon S3 buckets to discover personally identifiable information (Pll). The company stores the Pll data in the us-east-I Region and us-west-2 Region.
Which solution will meet these requirements with the LEAST operational overhead?

A. Configure Amazon Macie in each Regio
B. Create a job to analyze the data that is in Amazon S3_
C. Configure AWS Security Hub for all Region
D. Create an AWS Config rule to analyze the data that is in Amazon S3_
E. Configure Amazon Inspector to analyze the data that IS in Amazon S3.
F. Configure Amazon GuardDuty to analyze the data that is in Amazon S3.

Answer: A

Explanation:
it allows the solutions architect to review the S3 buckets to discover personally identifiable information (Pll) with the least operational overhead. Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data in AWS. Amazon Macie can analyze data in S3 buckets across multiple regions and provide insights into the type, location, and level of sensitivity of the data. References:
? Amazon Macie
? Analyzing data with Amazon Macie

NEW QUESTION 139
- (Topic 4)
A company is subscribed to the AWS Business Support plan. Compliance rules require the company to check on AWS infrastructure health before deployments can proceed. The company needs a programmatic and automated way to check on infrastructure health at the beginning of new deployments.
Which solution will meet these requirements?

A. Use the AWS Trusted Advisor API at the start of each deploymen
B. Pause all new deployments if the API returns any issues.
C. Use the AWS Health API at the start of each deploymen
D. Pause all new deployments if the API returns any issues.
E. Query the AWS Support API at the start of each deploymen
F. Pause all new deployments if the API returns any open issues.
G. Send an API call to each workload ahead of deploymen
H. Pause the deployments if the API call fails.

Answer: B

Explanation:
The AWS Health API provides programmatic access to the AWS Health information that is presented in the AWS Personal Health Dashboard. You can use the API operations to get information about AWS Health events that affect your AWS services and resources. You can also use the API to enable or disable health-based insights for your organization. You can use the AWS Health API at the start of each deployment to check on AWS infrastructure health and pause all new deployments if the API returns any issues. References: https://docs.aws.amazon.com/health/latest/APIReference/Welcome.html

NEW QUESTION 140
- (Topic 4)
A company copies 200 TB of data from a recent ocean survey onto AWS Snowball Edge Storage Optimized devices. The company has a high performance computing (HPC) cluster that is hosted on AWS to look for oil and gas deposits. A solutions architect must provide the cluster with consistent sub-millisecond latency and high-throughput access to the data on the Snowball Edge Storage Optimized devices. The company is sending the devices back to AWS.
Which solution will meet these requirements?

A. Create an Amazon S3 bucke
B. Import the data into the S3 bucke
C. Configure an AWS Storage Gateway file gateway to use the S3 bucke
D. Access the file gateway from the HPC cluster instances.
E. Create an Amazon S3 bucke
F. Import the data into the S3 bucke
G. Configure an Amazon FSx for Lustre file system, and integrate it with the S3 bucke
H. Access the FSx for Lustre file system from the HPC cluster instances.
I. Create an Amazon S3 bucket and an Amazon Elastic File System (Amazon EFS) file syste
J. Import the data into the S3 bucke
K. Copy the data from the S3 bucket to the EFS file syste
L. Access the EFS file system from the HPC cluster instances.
M. Create an Amazon FSx for Lustre file syste
N. Import the data directly into the FSx for Lustre file syste
O. Access the FSx for Lustre file system from the HPC cluster instances.

Answer: B

Explanation:
To provide the HPC cluster with consistent sub-millisecond latency and high- throughput access to the data on the Snowball Edge Storage Optimized devices, a solutions architect should configure an Amazon FSx for Lustre file system, and integrate it with an Amazon S3 bucket. This solution has the following benefits:
? It allows the HPC cluster to access the data on the Snowball Edge devices using a
POSIX-compliant file system that is optimized for fast processing of large datasets1.
? It enables the data to be imported from the Snowball Edge devices into the S3
bucket using the AWS Snow Family Console or the AWS CLI2. The data can then be accessed from the FSx for Lustre file system using the S3 integration feature3.
? It supports high availability and durability of the data, as the FSx for Lustre file
system can automatically copy the data to and from the S3 bucket3. The data can also be accessed from other AWS services or applications using the S3 API4.
References:
? 1: https://aws.amazon.com/fsx/lustre/
? 2: https://docs.aws.amazon.com/snowball/latest/developer-guide/using- adapter.html
? 3: https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data- repo.html
? 4: https://docs.aws.amazon.com/fsx/latest/LustreGuide/export-data-repo.html

NEW QUESTION 145
- (Topic 4)
A company's data platform uses an Amazon Aurora MySQL database. The database has multiple read replicas and multiple DB instances across different Availability Zones. Users have recently reported errors from the database that indicate that there are too many connections. The company wants to reduce the failover time by 20% when a read replica is promoted to primary writer.
Which solution will meet this requirement?

A. Switch from Aurora to Amazon RDS with Multi-AZ cluster deployment.
B. Use Amazon RDS Proxy in front of the Aurora database.
C. Switch to Amazon DynamoDB with DynamoDB Accelerator (DAX) for read connections.
D. Switch to Amazon Redshift with relocation capability.

Answer: B

Explanation:
Amazon RDS Proxy is a service that provides a fully managed, highly available database proxy for Amazon RDS and Aurora databases. It allows you to pool and share database connections, reduce database load, and improve application scalability and availability.
By using Amazon RDS Proxy in front of your Aurora database, you can achieve the following benefits:
? You can reduce the number of connections to your database and avoid errors that
indicate that there are too many connections. Amazon RDS Proxy handles the connection management and multiplexing for you, so you can use fewer database connections and resources.
? You can reduce the failover time by 20% when a read replica is promoted to
primary writer. Amazon RDS Proxy automatically detects failures and routes traffic to the new primary instance without requiring changes to your application code or configuration. According to a benchmark test, using Amazon RDS Proxy reduced the failover time from 66 seconds to 53 seconds, which is a 20% improvement.
? You can improve the security and compliance of your database access. Amazon
RDS Proxy integrates with AWS Secrets Manager and AWS Identity and Access Management (IAM) to enable secure and granular authentication and authorization for your database connections.

NEW QUESTION 147
.....

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