AWS-Solution-Architect-Associate Dumps

Amazon AWS Certified Solutions Architect - Associate

https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html

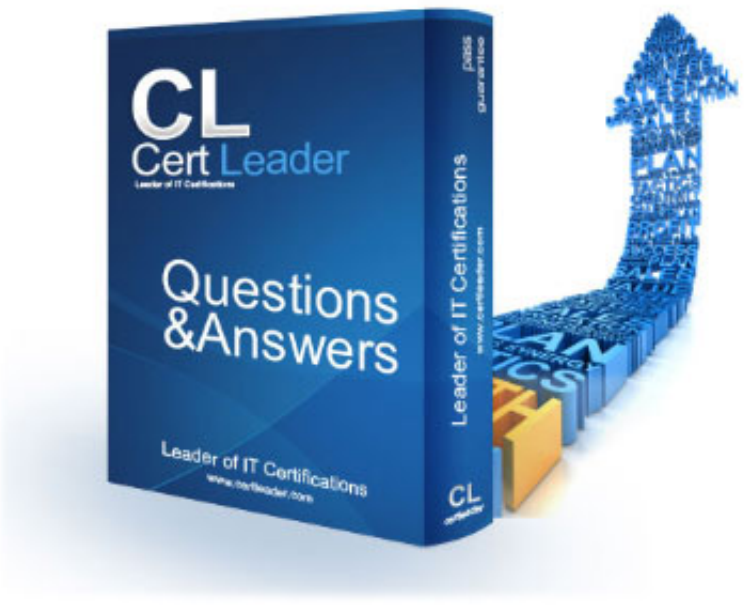

NEW QUESTION 1
- (Topic 4)
A solutions architect needs to design the architecture for an application that a vendor provides as a Docker container image The container needs 50 GB of storage available for temporary files The infrastructure must be serverless.
Which solution meets these requirements with the LEAST operational overhead?

A. Create an AWS Lambda function that uses the Docker container image with an Amazon S3 mounted volume that has more than 50 GB of space
B. Create an AWS Lambda function that uses the Docker container image with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space
C. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the AWS Fargate launch type Create a task definition for the container image with an Amazon Elastic File System (Amazon EFS) volum
D. Create a service with that task definition.
E. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the Amazon EC2 launch type with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space Create a task definition for the container imag
F. Create a service with that task definition.

Answer: C

Explanation:
The AWS Fargate launch type is a serverless way to run containers on Amazon ECS,
without having to manage any underlying infrastructure. You only pay for the resources required to run your containers, and AWS handles the provisioning, scaling, and security of the cluster. Amazon EFS is a fully managed, elastic, and scalable file system that can be mounted to multiple containers, and provides high availability and durability. By using AWS Fargate and Amazon EFS, you can run your Docker container image with 50 GB of storage available for temporary files, with the least operational overhead. This solution meets the requirements of the question.
References:
? AWS Fargate
? Amazon Elastic File System
? Using Amazon EFS file systems with Amazon ECS

NEW QUESTION 2
- (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 3
- (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 4
- (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 5
- (Topic 4)
A company has an application that uses Docker containers in its local data center The application runs on a container host that stores persistent data in a volume on the host. The container instances use the stored persistent data.
The company wants to move the application to a fully managed service because the company does not want to manage any servers or storage infrastructure.
Which solution will meet these requirements?

A. Use Amazon Elastic Kubernetes Service (Amazon EKS) with self-managed node
B. Create an Amazon Elastic Block Store (Amazon EBS) volume attached to an Amazon EC2 instanc
C. Use the EBS volume as a persistent volume mounted in the containers.
D. Use Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch typ
E. Create an Amazon Elastic File System (Amazon EFS) volum
F. Add the EFS volumeas a persistent storage volume mounted in the containers.
G. Use Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch typ
H. Create an Amazon S3 bucke
I. Map the S3 bucket as a persistent storage volume mounted in the containers.
J. Use Amazon Elastic Container Service (Amazon ECS) with an Amazon EC2 launch typ
K. Create an Amazon Elastic File System (Amazon EFS) volum
L. Add the EFS volume as a persistent storage volume mounted in the containers.

Answer: B

Explanation:
This solution meets the requirements because it allows the company to move the application to a fully managed service without managing any servers or storage infrastructure. AWS Fargate is a serverless compute engine for containers that runs the Amazon ECS tasks. With Fargate, the company does not need to provision, configure, or scale clusters of virtual machines to run containers. Amazon EFS is a fully managed file system that can be accessed by multiple containers concurrently. With EFS, the company does not need to provision and manage storage capacity. EFS provides a simple interface to create and configure file systems quickly and easily. The company can use the EFS volume as a persistent storage volume mounted in the containers to store the persistent data. The company can also use the EFS mount helper to simplify the mounting process. References: Amazon ECS on AWS Fargate, Using Amazon EFS file systems with Amazon ECS, Amazon EFS mount helper.

NEW QUESTION 6
- (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 7
- (Topic 4)
A company offers a food delivery service that is growing rapidly. Because of the growth, the company's order processing system is experiencing scaling problems during peak traffic hours. The current architecture includes the following:
• A group of Amazon EC2 instances that run in an Amazon EC2 Auto Scaling group to collect orders from the application
• Another group of EC2 instances that run in an Amazon EC2 Auto Scaling group to fulfill orders
The order collection process occurs quickly, but the order fulfillment process can take longer. Data must not be lost because of a scaling event.
A solutions architect must ensure that the order collection process and the order fulfillment process can both scale properly during peak traffic hours. The solution must optimize
utilization of the company's AWS resources. Which solution meets these requirements?

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

Answer: D

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

NEW QUESTION 8
- (Topic 4)
A company has an on-premises server that uses an Oracle database to process and store customer information The company wants to use an AWS database service to achieve higher availability and to improve application performance. The company also wants to offload reporting from its primary database system. Which solution will meet these requirements in the MOST operationally efficient way?

A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions Point the reporting functions toward a separate DB instance from the primary DB instance.
B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database Create a read replica in the same zone as the primary DB instanc
C. Direct the reporting functions to the read replica.
D. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database Direct the reporting functions to use the reader instance in the cluster deployment
E. Use Amazon RDS deployed in a Multi-AZ instance deployment to create an Amazon Aurora databas
F. Direct the reporting functions to the reader instances.

Answer: D

Explanation:
Amazon Aurora is a fully managed relational database that is compatible with MySQL and PostgreSQL. It provides up to five times better performance than MySQL and
up to three times better performance than PostgreSQL. It also provides high availability and durability by replicating data across multiple Availability Zones and continuously backing up data to Amazon S31. By using Amazon RDS deployed in a Multi-AZ instance deployment
to create an Amazon Aurora database, the solution can achieve higher availability and improve application performance.
Amazon Aurora supports read replicas, which are separate instances that share the same underlying storage as the primary instance. Read replicas can be used to offload read-only queries from the primary instance and improve performance. Read replicas can also be used for reporting functions2. By directing the reporting functions to the reader instances, the solution can offload reporting from its primary database system.
* A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions Point the reporting functions toward a separate DB instance from the pri-mary DB instance. This solution will not meet the requirement of using an AWS database service, as AWS DMS is a service that helps users migrate databases to AWS, not a database service itself. It also involves creating multiple DB instances in different Regions, which may increase complexity and cost.
* B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database Create a read replica in the same zone as the primary DB instance. Direct the reporting functions to the read replica. This solution will not meet the requirement of achieving higher availability, as a Single-AZ deployment does not provide failover protection in case of an Availability Zone outage. It also involves using Oracle as the database engine, which may not provide better performance than Aurora.
* C. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database Di-rect the reporting functions to use the reader instance in the cluster deployment. This solution will not meet the requirement of improving application performance, as Oracle may not provide better performance than Aurora. It also involves using a cluster deployment, which is only supported for Aurora, not for Oracle. Reference URL: https://aws.amazon.com/rds/aurora/

NEW QUESTION 9
- (Topic 4)
A manufacturing company runs its report generation application on AWS. The application generates each report in about 20 minutes. The application is built as a monolith that runs on a single Amazon EC2 instance. The application requires frequent updates to its tightly coupled modules. The application becomes complex to maintain as the company adds new features.
Each time the company patches a software module, the application experiences downtime. Report generation must restart from the beginning after any interruptions. The company wants to redesign the application so that the application can be flexible, scalable, and gradually improved. The company wants to

minimize application downtime.
Which solution will meet these requirements?

A. Run the application on AWS Lambda as a single function with maximum provisioned concurrency.
B. Run the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy.
C. Run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling.
D. Run the application on AWS Elastic Beanstalk as a single application environment with an all-at-once deployment strategy.

Answer: C

Explanation:
The solution that will meet the requirements is to run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling. This solution will allow the application to be flexible, scalable, and gradually improved, as well as minimize application downtime. By breaking down the monolithic application into microservices, the company can decouple the modules and update them independently, without affecting the whole application. By running the microservices on Amazon ECS, the company can leverage the benefits of containerization, such as portability, efficiency, and isolation. By enabling service auto scaling, the company can adjust the number of containers running for each microservice based on demand, ensuring optimal performance and cost. Amazon ECS also supports various deployment strategies, such as rolling update or blue/green deployment, that can reduce or eliminate downtime during updates.
The other solutions are not as effective as the first one because they either do not meet the requirements or introduce new challenges. Running the application on AWS Lambda as a single function with maximum provisioned concurrency will not meet the requirements, as it will not break down the monolith into microservices, nor will it reduce the complexity of maintenance. Lambda functions are also limited by execution time (15 minutes), memory size (10 GB), and concurrency quotas, which may not be sufficient for the report generation application. Running the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy will not meet the requirements, as it will introduce the risk of interruptions due to spot price fluctuations. Spot Instances are not guaranteed to be available or stable, and may be reclaimed by AWS at any time with a two-minute warning. This may cause report generation to fail or restart from scratch. Running the application on AWS Elastic Beanstalk as a single application environment with an all-at- once deployment strategy will not meet the requirements, as it will not break down the monolith into microservices, nor will it minimize application downtime. The all-at-once deployment strategy will deploy updates to all instances simultaneously, causing a brief outage for the application.
References:
? Amazon Elastic Container Service
? Microservices on AWS
? Service Auto Scaling - Amazon Elastic Container Service
? AWS Lambda
? Amazon EC2 Spot Instances
? [AWS Elastic Beanstalk]

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 10
- (Topic 4)
A company uses Amazon EC2 instances to host its internal systems. As part of a deployment operation, an administrator tries to use the AWS CLI to terminate an EC2 instance. However, the administrator receives a 403 (Access Denied) error message.
The administrator is using an IAM role that has the following IAM policy attached:

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

What is the cause of the unsuccessful request?

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

Answer: D

NEW QUESTION 13
- (Topic 4)
A company wants to migrate its three-tier application from on premises to AWS. The web tier and the application tier are running on third-party virtual machines (VMs). The database tier is running on MySQL.
The company needs to migrate the application by making the fewest possible changes to the architecture. The company also needs a database solution that can restore data to a specific point in time.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

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

NEW QUESTION 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 20
- (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 is building an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for its workloads. All secrets that are stored in Amazon EKS must be encrypted in the Kubernetes etcd key-value store.
Which solution will meet these requirements?

A. Create a new AWS Key Management Service (AWS KMS) key Use AWS Secrets Manager to manage rotate, and store all secrets in Amazon EKS.
B. Create a new AWS Key Management Service (AWS KMS) key Enable Amazon EKS KMS secrets encryption on the Amazon EKS cluster.
C. Create the Amazon EKS cluster with default options Use the Amazon Elastic Block Store (Amazon EBS) Container Storage Interface (CSI) driver as an add-on.
D. Create a new AWS Key Management Service (AWS KMS) key with the ahas/aws/ebs alias Enable default Amazon Elastic Block Store (Amazon EBS) volume encryption for the account.

Answer: B

Explanation:
This option is the most secure and simple way to encrypt the secrets that are stored in Amazon EKS. AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys that can be used to encrypt your data. Amazon EKS KMS secrets encryption is a feature that enables you to use a KMS key to encrypt the secrets that are stored in the Kubernetes etcd key-value store. This provides an additional layer of protection for your sensitive data, such as passwords, tokens, and keys. You can create a new KMS key or use an existing one, and then enable the Amazon EKS KMS secrets encryption on the Amazon EKS cluster. You can also use IAM policies to control who can access or use the KMS key.
Option A is not correct because using AWS Secrets Manager to manage, rotate, and store all secrets in Amazon EKS is not necessary or efficient. AWS Secrets Manager is a service that helps you securely store, retrieve, and rotate your secrets, such as database credentials, API keys, and passwords. You can use it to manage secrets that are used by your applications or services outside of Amazon EKS, but it is not designed to encrypt the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using AWS Secrets Manager would incur additional costs and complexity, and it would not leverage the native Kubernetes secrets management capabilities.
Option C is not correct because using the Amazon EBS Container Storage Interface (CSI) driver as an add-on does not encrypt the secrets that are stored in Amazon EKS. The Amazon EBS CSI driver is a plugin that allows you to use Amazon EBS volumes as persistent storage for your Kubernetes pods. It is useful for providing durable and scalable storage for your applications, but it does not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the Amazon EBS CSI driver would require additional configuration and resources, and it would not provide the same level of security as using a KMS key.
Option D is not correct because creating a new AWS KMS key with the alias aws/ebs and enabling default Amazon EBS volume encryption for the account does not encrypt the secrets that are stored in Amazon EKS. The alias aws/ebs is a reserved alias that is used by AWS to create a default KMS key for your account. This key is used to encrypt the Amazon EBS volumes that are created in your account, unless you specify a different KMS key. Enabling default Amazon EBS volume encryption for the account is a setting that ensures that all new Amazon EBS volumes are encrypted by default. However, these features do not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the default KMS key or the default encryption setting would not provide the same level of control and security as using a custom KMS key and enabling the Amazon EKS KMS secrets encryption feature. References:
? Encrypting secrets used in Amazon EKS
? What Is AWS Key Management Service?
? What Is AWS Secrets Manager?
? Amazon EBS CSI driver
? Encryption at rest

NEW QUESTION 26
- (Topic 4)
A city has deployed a web application running on Amazon EC2 instances behind an Application Load Balancer (ALB). The application's users have reported sporadic performance, which appears to be related to DDoS attacks originating from random IP addresses. The city needs a solution that requires minimal configuration changes and provides an audit trail for the DDoS sources.
Which solution meets these requirements?

A. Enable an AWS WAF web ACL on the ALB, and configure rules to block traffic from unknown sources.
B. Subscribe to Amazon Inspecto
C. Engage the AWS DDoS Response Team (DRT) to integrate mitigating controls into the service.
D. Subscribe to AWS Shield Advance
E. Engage the AWS DDoS Response Team (DRT) to integrate mitigating controls into the service.
F. Create an Amazon CloudFront distribution for the application, and set the ALB as the origi
G. Enable an AWS WAF web ACL on the distribution, and configure rules to block traffic from unknown sources.

Answer: C

Explanation:
To protect the web application from DDoS attacks originating from random IP addresses, a solutions architect should subscribe to AWS Shield Advanced and engage the AWS DDoS Response Team (DRT) to integrate mitigating controls into the service. AWS Shield Advanced is a managed service that provides protection against large and sophisticated DDoS attacks, with access to 24/7 support and response from the DRT. The DRT can help the city configure proactive and reactive safeguards, such as AWS WAF rules, rate-based rules, and network ACLs, to block malicious traffic and improve the application’s resilience. The

service also provides an audit trail for the DDoS sources through detailed attack reports and Amazon CloudWatch metrics.

NEW QUESTION 31
- (Topic 4)
A company runs a web application that is deployed on Amazon EC2 instances in the private subnet of a VPC. An Application Load Balancer (ALB) that extends across the public subnets directs web traffic to the EC2 instances. The company wants to implement new security measures to restrict inbound traffic from the ALB to the EC2 instances while preventing access from any other source inside or outside the private subnet of the EC2 instances.
Which solution will meet these requirements?

A. Configure a route in a route table to direct traffic from the internet to the private IP addresses of the EC2 instances.
B. Configure the security group for the EC2 instances to only allow traffic that comes from the security group for the ALB.
C. Move the EC2 instances into the public subne
D. Give the EC2 instances a set of Elastic IP addresses.
E. Configure the security group for the ALB to allow any TCP traffic on any port.

Answer: B

Explanation:
To restrict inbound traffic from the ALB to the EC2 instances, the security group for the EC2 instances should only allow traffic that comes from the security group for the ALB. This way, the EC2 instances can only receive requests from the ALB and not from any other source inside or outside the private subnet.
References:
? Security Groups for Your Application Load Balancers
? Security Groups for Your VPC

NEW QUESTION 33
- (Topic 4)
A company has deployed a multiplayer game for mobile devices. The game requires live
location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations.
The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly.
What should a solutions architect do to improve the performance of the data tier?

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

Answer: D

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

NEW QUESTION 34
- (Topic 4)
A company needs to integrate with a third-party data feed. The data feed sends a webhook to notify an external service when new data is ready for consumption A developer wrote an AWS Lambfe function to retrieve data when the company receives a webhook callback The developer must make the Lambda function available for the third party to call.
Which solution will meet these requirements with the MOST operational efficiency?

A. Create a function URL for the Lambda functio
B. Provide the Lambda function URL to the third party for the webhook.
C. Deploy an Application Load Balancer (ALB) in front of the Lambda functio
D. Provide the ALB URL to the third party for the webhook
E. Create an Amazon Simple Notification Service (Amazon SNS) topi
F. Attach the topic to the Lambda functio
G. Provide the public hostname of the SNS topic to the third party for the webhook.
H. Create an Amazon Simple Queue Service (Amazon SQS) queu
I. Attach the queue to the Lambda functio
J. Provide the public hostname of the SQS queue to the third party forthe webhook.

Answer: A

Explanation:
A function URL is a unique identifier for a Lambda function that can be used to invoke the function over HTTPS. It is composed of the API endpoint of the AWS Region where the function is deployed, and the name or ARN of the function1. By creating a function URL for the Lambda function, the solution can make the Lambda function available for the third party to call with the most operational efficiency.
* B. Deploy an Application Load Balancer (ALB) in front of the Lambda function. Provide the ALB URL to the third party for the webhook. This solution will not meet the requirement of the most operational efficiency, as it involves creating and managing an additional resource (ALB) that is not necessary for invoking a Lambda function over HTTPS2.
* C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Attach the topic to the Lambda function. Provide the public hostname of the SNS topic to the third party for the webhook. This solution will not work, as Amazon SNS topics do not have public hostnames that can be used as webhooks. SNS topics are used to publish messages to subscribers, not to receive messages from external sources3.
* D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Attach the queue to the Lamb-da function. Provide the public hostname of the SQS queue to the third party for the webhook. This solution will not work, as Amazon SQS queues do not have public hostnames that can be used as webhooks. SQS queues are used to send, store, and receive messages between AWS services, not to receive messages from external sources. Reference URL: https://docs.aws.amazon.com/lambda/latest/dg/lambda-api-permissions- ref.html

NEW QUESTION 36
- (Topic 4)
A company needs to minimize the cost of its 1 Gbps AWS Direct Connect connection. The company's average connection utilization is less than 10%. A solutions architect must recommend a solution that will reduce the cost without compromising security.
Which solution will meet these requirements?

A. Set up a new 1 Gbps Direct Connect connectio
B. Share the connection with another AWS account.
C. Set up a new 200 Mbps Direct Connect connection in the AWS Management Console.
D. Contact an AWS Direct Connect Partner to order a 1 Gbps connectio
E. Share the connection with another AWS account.
F. Contact an AWS Direct Connect Partner to order a 200 Mbps hosted connection for an existing AWS account.

Answer: D

Explanation:
company need to setup a cheaper connection (200 M) but B is incorrect because you can only order port speeds of 1, 10, or 100 Gbps for more flexibility you can go with hosted connection, You can order port speeds between 50 Mbps and 10 Gbps. https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct- connect.html

NEW QUESTION 37
- (Topic 4)
A company runs multiple workloads in its on-premises data center. The company's data center cannot scale fast enough to meet the company's expanding business needs. The company wants to collect usage and configuration data about the on-premises servers and workloads to plan a migration to AWS.
Which solution will meet these requirements?

A. Set the home AWS Region in AWS Migration Hu
B. Use AWS Systems Manager to collect data about the on-premises servers.
C. Set the home AWS Region in AWS Migration Hu
D. Use AWS Application Discovery Service to collect data about the on-premises servers.
E. Use the AWS Schema Conversion Tool (AWS SCT) to create the relevant template
F. Use AWS Trusted Advisor to collect data about the on-premises servers.
G. Use the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates.Use AWS Database Migration Service (AWS DMS) to collect data about the on-premises servers.

Answer: B

Explanation:
The most suitable solution for the company’s requirements is to set the home AWS Region in AWS Migration Hub and use AWS Application Discovery Service to collect data about the on-premises servers. This solution will enable the company to gather usage and configuration data of its on-premises servers and workloads, and plan a migration to AWS.
AWS Migration Hub is a service that simplifies and accelerates migration tracking by aggregating migration status information into a single console. Users can view the discovered servers, group them into applications, and track the migration status of each application from the Migration Hub console in their home Region. The home Region is the AWS Region where users store their migration data, regardless of which Regions they migrate into1.
AWS Application Discovery Service is a service that helps users plan their migration to AWS by collecting usage and configuration data about their on-premises servers and databases. Application Discovery Service is integrated with AWS Migration Hub and supports two methods of performing discovery: agentless discovery and agent-based discovery. Agentless discovery can be performed by deploying the Application Discovery Service Agentless Collector through VMware vCenter, which collects static configuration data and utilization data for virtual machines (VMs) and databases. Agent-based discovery can be performed by deploying the AWS Application Discovery Agent on each of the VMs and physical servers, which collects static configuration data, detailed time-series system-performance information, inbound and outbound network connections, and processes that are running2.
The other options are not correct because they do not meet the requirements or are not relevant for the use case. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Trusted Advisor to collect data about the on- premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another, such as from Oracle to PostgreSQL3. AWS Trusted Advisor is a service that provides best practice recommendations for cost optimization, performance, security, fault tolerance, and service limits4. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Database Migration Service (AWS DMS) to collect data about the on-premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. As mentioned above, AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another. AWS DMS is a service that helps users migrate relational databases, non-relational databases, and other types of data stores to
AWS with minimal downtime5. References:
? Home Region - AWS Migration Hub
? What is AWS Application Discovery Service? - AWS Application Discovery Service
? AWS Schema Conversion Tool - Amazon Web Services
? What Is Trusted Advisor? - Trusted Advisor
? What Is AWS Database Migration Service? - AWS Database Migration Service

NEW QUESTION 42
- (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 44
- (Topic 4)
A company hosts an application used to upload files to an Amazon S3 bucket Once uploaded, the files are processed to extract metadata which takes less than 5 seconds The volume and frequency of the uploads varies from a few files each hour to hundreds of concurrent uploads The company has asked a solutions architect to design a cost-effective architecture that will meet these requirements.
What should the solutions architect recommend?

A. Configure AWS CloudTrail trails to tog S3 API calls Use AWS AppSync to process the files.
B. Configure an object-created event notification within the S3 bucket to invoke an AWS Lambda function to process the files.
C. Configure Amazon Kinesis Data Streams to process and send data to Amazon S3. Invoke an AWS Lambda function to process the files.
D. Configure an Amazon Simple Notification Service (Amazon SNS) topic to process the files uploaded to Amazon S3 Invoke an AWS Lambda function to process the files.

Answer: B

Explanation:
This option is the most cost-effective and scalable way to process the files uploaded to S3. AWS CloudTrail is used to log API calls, not to trigger actions based on them. AWS AppSync is a service for building GraphQL APIs, not for processing files. Amazon Kinesis Data Streams is used to ingest and process streaming data, not to send data to S3. Amazon SNS is a pub/sub service that can be used to notify subscribers of events, not to process files. References:
? Using AWS Lambda with Amazon S3
? AWS CloudTrail FAQs
? What Is AWS AppSync?
? [What Is Amazon Kinesis Data Streams?]
? [What Is Amazon Simple Notification Service?]

NEW QUESTION 46
- (Topic 4)
A company is deploying an application that processes streaming data in near-real time The company plans to use Amazon EC2 instances for the workload The network architecture must be configurable to provide the lowest possible latency between nodes
Which combination of network solutions will meet these requirements? (Select TWO)

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

Answer: AC

Explanation:
These options are the most suitable ways to configure the network architecture to provide the lowest possible latency between nodes. Option A enables and configures enhanced networking on each EC2 instance, which is a feature that improves the network performance of the instance by providing higher bandwidth,

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

NEW QUESTION 51
- (Topic 4)
A company runs analytics software on Amazon EC2 instances The software accepts job requests from users to process data that has been uploaded to Amazon S3 Users report that some submitted data is not being processed Amazon CloudWatch reveals that the EC2 instances have a consistent CPU utilization at or near 100% The company wants to improve system performance and scale the system based on user load.
What should a solutions architect do to meet these requirements?

A. Create a copy of the instance Place all instances behind an Application Load Balancer
B. Create an S3 VPC endpoint for Amazon S3 Update the software to reference the endpoint
C. Stop the EC2 instance
D. Modify the instance type to one with a more powerful CPU and more memor
E. Restart the instances.
F. Route incoming requests to Amazon Simple Queue Service (Amazon SQS) Configure an EC2 Auto Scaling group based on queue size Update the software to read from thequeue.

Answer: D

Explanation:
This option is the best solution because it allows the company to decouple the analytics software from the user requests and scale the EC2 instances dynamically based on the demand. By using Amazon SQS, the company can create a queue that stores the user requests and acts as a buffer between the users and the analytics software. This way, the software can process the requests at its own pace without losing any data or overloading the EC2 instances. By using EC2 Auto Scaling, the company can create an Auto Scaling group that launches or terminates EC2 instances automatically based on the size of the queue. This way, the company can ensure that there are enough instances to handle the load and optimize the cost and performance of the system. By updating the software to read from the queue, the company can enable the analytics software to consume the requests from the queue and process the data from Amazon S3.
* A. Create a copy of the instance Place all instances behind an Application Load Balancer. This option is not optimal because it does not address the root cause of the problem, which is the high CPU utilization of the EC2 instances. An Application Load Balancer can distribute the incoming traffic across multiple instances, but it cannot scale the instances based on the load or reduce the processing time of the analytics software. Moreover, this option can incur additional costs for the load balancer and the extra instances.
* B. Create an S3 VPC endpoint for Amazon S3 Update the software to reference the endpoint. This option is not effective because it does not solve the issue of the high CPU utilization of the EC2 instances. An S3 VPC endpoint can enable the EC2 instances to access Amazon S3 without going through the internet, which can improve the network performance and security. However, it cannot reduce the processing time of the analytics software or scale the instances based on the load.
* C. Stop the EC2 instances. Modify the instance type to one with a more powerful CPU and more memory. Restart the instances. This option is not scalable because it does not account for the variability of the user load. Changing the instance type to a more powerful one can improve the performance of the analytics software, but it cannot adjust the number of instances based on the demand. Moreover, this option can increase the cost of the system and cause downtime during the instance modification.
References:
? 1 Using Amazon SQS queues with Amazon EC2 Auto Scaling - Amazon EC2 Auto Scaling
? 2 Tutorial: Set up a scaled and load-balanced application - Amazon EC2 Auto Scaling
? 3 Amazon EC2 Auto Scaling FAQs

NEW QUESTION 56
- (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 58
- (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 60
- (Topic 4)
A company is running a photo hosting service in the us-east-1 Region. The service enables users across multiple countries to upload and view photos. Some photos are heavily viewed for months, and others are viewed for less than a week. The application allows uploads of up to 20 MB for each photo. The service uses the photo metadata to determine which photos to display to each user.
Which solution provides the appropriate user access MOST cost-effectively?

A. Store the photos in Amazon DynamoD
B. Turn on DynamoDB Accelerator (DAX) to cache frequently viewed items.
C. Store the photos in the Amazon S3 Intelligent-Tiering storage clas
D. Store the photo metadata and its S3 location in DynamoDB.
E. Store the photos in the Amazon S3 Standard storage clas
F. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Standard-Infrequent Access (S3 Standard-IA) storage clas
G. Use the object tags to keep track of metadata.
H. Store the photos in the Amazon S3 Glacier storage clas
I. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Glacier Deep Archive storage clas
J. Store the photo metadata and its S3 location in Amazon OpenSearch Service.

Answer: B

Explanation:
This solution provides the appropriate user access most cost-effectively
because it uses the Amazon S3 Intelligent-Tiering storage class, which automatically optimizes storage costs by moving data to the most cost-effective access tier when access patterns change, without performance impact or operational overhead1. This storage class is ideal for data with unknown, changing, or unpredictable access patterns, such as photos that are heavily viewed for months or less than a week. By storing the photo metadata and its S3 location in DynamoDB, the application can quickly query and retrieve the relevant photos for each user. DynamoDB is a fast, scalable, and fully managed NoSQL database service that supports key-value and document data models2.
References: 1: Amazon S3 Intelligent-Tiering Storage Class | AWS3, Overview section2: Amazon DynamoDB - NoSQL Cloud Database Service4, Overview section.

NEW QUESTION 63
- (Topic 4)
A social media company wants to allow its users to upload images in an application that is hosted in the AWS Cloud. The company needs a solution that automatically resizes the images so that the images can be displayed on multiple device types. The application experiences unpredictable traffic patterns throughout the day. The company is seeking a highly available solution that maximizes scalability.
What should a solutions architect do to meet these requirements?

A. Create a static website hosted in Amazon S3 that invokes AWS Lambda functions to resize the images and store the images in an Amazon S3 bucket.
B. Create a static website hosted in Amazon CloudFront that invokes AWS Step Functions to resize the images and store the images in an Amazon RDS database.
C. Create a dynamic website hosted on a web server that runs on an Amazon EC2 instance Configure a process that runs on the EC2 instance to resize the images and store the images in an Amazon S3 bucket.
D. Create a dynamic website hosted on an automatically scaling Amazon Elastic Container Service (Amazon ECS) cluster that creates a resize job in Amazon Simple Queue Service (Amazon SQS). Set up an image-resizing program that runs on an Amazon EC2 instance to process the resize jobs

Answer: A

Explanation:
By using Amazon S3 and AWS Lambda together, you can create a serverless architecture that provides highly scalable and available image resizing capabilities. Here's how the solution would work: Set up an Amazon S3 bucket to store the original images uploaded by users. Configure an event trigger on the S3 bucket to invoke an AWS Lambda function whenever a new image is uploaded. The Lambda function can be designed to retrieve the uploaded image, perform the

necessary resizing operations based on device requirements, and store the resized images back in the S3 bucket or a different bucket designated for resized images. Configure the Amazon S3 bucket to make the resized images publicly accessible for serving to users.

NEW QUESTION 68
- (Topic 4)
A company has a large workload that runs every Friday evening. The workload runs on Amazon EC2 instances that are in two Availability Zones in the us-east-1 Region. Normally, the company must run no more than two instances at all times. However, the company wants to scale up to six instances each Friday to handle a regularly repeating increased workload.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: B

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

NEW QUESTION 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 73
- (Topic 4)
A company uses AWS Organizations. The company wants to operate some of its AWS accounts with different budgets. The company wants to receive alerts and automatically prevent provisioning of additional resources on AWS accounts when the allocated budget threshold is met during a specific period.
Which combination of solutions will meet these requirements? (Select THREE.)

A. Use AWS Budgets to create a budge
B. Set the budget amount under the Cost and Usage Reports section of the required AWS accounts.
C. Use AWS Budgets to create a budge
D. Set the budget amount under the Billing dashboards of the required AWS accounts.
E. Create an 1AM user for AWS Budgets to run budget actions with the required permissions.
F. Create an 1AM role for AWS Budgets to run budget actions with the required permissions.
G. Add an alert to notify the company when each account meets its budget threshol
H. Add a budget action that selects the 1AM identity created with the appropriate config rule to prevent provisioning of additional resources.
I. Add an alert to notify the company when each account meets its budget threshol
J. Add a budget action that selects the 1AM identity created with the appropriate service control policy (SCP) to prevent provisioning of additional resources.

Answer: BDF

Explanation:
To use AWS Budgets to create and manage budgets for different AWS accounts, the company needs to do the following steps:
? Use AWS Budgets to create a budget for each AWS account that needs a different
budget amount. The budget can be based on cost or usage metrics, and can have different time periods, filters, and thresholds. The company can set the budget amount under the Billing dashboards of the required AWS accounts1.
? Create an IAM role for AWS Budgets to run budget actions with the required
permissions. A budget action is a response that AWS Budgets initiates when a
budget exceeds a specified threshold. The IAM role allows AWS Budgets to perform actions on behalf of the company, such as applying an IAM policy or a service control policy (SCP) to restrict the provisioning of additional resources2.
? Add an alert to notify the company when each account meets its budget threshold.
The alert can be sent via email or Amazon SNS. The company can also add a budget action that selects the IAM role created and the appropriate SCP to prevent provisioning of additional resources. An SCP is a type of policy that can be applied to an AWS account or an organizational unit (OU) within AWS Organizations. An SCP can limit the actions that users and roles can perform in the account or OU3.

References:
? 4: https://aws.amazon.com/budgets/
? 1: https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets- create.html
? 2: https://docs.aws.amazon.com/cost-management/latest/userguide/budgets- controls.html
? 3:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policie s_scps.html

NEW QUESTION 78
- (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 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 86
- (Topic 4)
A company is running its production and nonproduction environment workloads in multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to design a solution that will prevent the modification of cost usage tags.
Which solution will meet these requirements?

A. Create a custom AWS Config rule to prevent tag modification except by authorized principals.
B. Create a custom trail in AWS CloudTrail to prevent tag modification
C. Create a service control policy (SCP) to prevent tag modification except by authonzed principals.
D. Create custom Amazon CloudWatch logs to prevent tag modification.

Answer: C

Explanation:
This solution meets the requirements because it uses SCPs to restrict the actions that can be performed on cost usage tags in the organization. SCPs are a type of organization policy that you can use to manage permissions in your organization. SCPs specify the maximum permissions for an organization, organizational

unit (OU), or account. You can use SCPs to enforce consistent tag policies across your organization and prevent unauthorized or accidental changes to your tags. You can also create exceptions for authorized principals, such as administrators or auditors, who need to modify tags for legitimate purposes.
References:
? Service control policies (SCPs) - AWS Organizations
? Tag policies - AWS Organizations

NEW QUESTION 91
- (Topic 4)
A company has applications hosted on Amazon EC2 instances with IPv6 addresses. The applications must initiate communications with other external applications using the internet.
However, the company's security policy states that any external service cannot initiate a connection to the EC2 instances.
What should a solutions architect recommend to resolve this issue?

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

Answer: D

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

NEW QUESTION 93
- (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 imagemanagement 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:
To resize images dynamically and serve appropriate formats to clients, a Lambda@Edge function with an external image management library can be used. Lambda@Edge allows running custom code at the edge locations of CloudFront, which can process the images on the fly and optimize them for different devices and browsers. An external image management library can provide various image manipulation and optimization features. References:
? Lambda@Edge
? Resizing Images with Amazon CloudFront & Lambda@Edge

NEW QUESTION 94
- (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 96
- (Topic 4)
A company runs applications on AWS that connect to the company's Amazon RDS database. The applications scale on weekends and at peak times of the year. The company wants to scale the database more effectively for its applications that connect to the database.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon DynamoDB with connection pooling with a target group configuration for the databas
B. Change the applications to use the DynamoDB endpoint.
C. Use Amazon RDS Proxy with a target group for the databas
D. Change the applications to use the RDS Proxy endpoint.
E. Use a custom proxy that runs on Amazon EC2 as an intermediary to the databas
F. Change the applications to use the custom proxy endpoint.
G. Use an AWS Lambda function to provide connection pooling with a target group configuration for the databas
H. Change the applications to use the Lambda function.

Answer: B

Explanation:
Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure1. RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability2. RDS Proxy also reduces failover times for Aurora and RDS databases by up to 66% and enables IAM authentication and Secrets Manager integration for database access1. RDS Proxy can be enabled for most applications with no code changes2.

NEW QUESTION 98
- (Topic 4)
A serverless application uses Amazon API Gateway. AWS Lambda, and Amazon DynamoDB. The Lambda function needs permissions to read and write to the DynamoDB table.
Which solution will give the Lambda function access to the DynamoDB table MOST securely?

A. Create an 1AM user with programmatic access to the Lambda functio
B. Attach a policy to the user that allows read and write access to the DynamoDB tabl
C. Store the access_key_id and secret_access_key parameters as part of the Lambda environment variable
D. Ensure that other AWS users do not have read and write access to the Lambda function configuration
E. Create an 1AM role that includes Lambda as a trusted servic
F. Attach a policy to the role that allows read and write access to the DynamoDB tabl
G. Update the configuration of the Lambda function to use the new role as the execution role.
H. Create an 1AM user with programmatic access to the Lambda functio
I. Attach a policy to the user that allows read and write access to the DynamoDB tabl
J. Store the access_key_id and secret_access_key parameters in AWS Systems Manager Parameter Store as secure string parameter
K. Update the Lambda function code to retrieve the secure string parameters before connecting to the DynamoDB table.
L. Create an 1AM role that includes DynamoDB as a trusted servic
M. Attach a policy to the role that allows read and write access from the Lambda functio
N. Update the code of the Lambda function to attach to the new role as an execution role.

Answer: B

Explanation:
Option B suggests creating an IAM role that includes Lambda as a trusted service, meaning the role is specifically designed for Lambda functions. The role should have a policy attached to it that grants the required read and write access to the DynamoDB table.

NEW QUESTION 102
- (Topic 4)
A company wants to run its payment application on AWS The application receives payment notifications from mobile devices Payment notifications require a basic validation before they are sent for further processing
The backend processing application is long running and requires compute and memory to be adjusted The company does not want to manage the infrastructure
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Amazon Simple Queue Service (Amazon SQS) queue Integrate the queue with an Amazon EventBndge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS) Anywhere Create a standalone cluster
B. Create an Amazon API Gateway API Integrate the API with anAWS Step Functions state machine to receive payment notifications from mobile devices Invoke the state machine to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Sen/ice (Amazon EKS). Configure an EKS cluster with self-managed nodes.
C. Create an Amazon Simple Queue Sen/ice (Amazon SQS) queue Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon EC2 Spot Instances Configure a Spot Fleet with a default allocation strategy.
D. Create an Amazon API Gateway API Integrate the API with AWS Lambda to receive payment notifications from mobile devices Invoke a Lambda function to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Container Service (Amazon

ECS). Configure Amazon ECS with an AWS Fargate launch type.

Answer: D

Explanation:
This option is the best solution because it allows the company to run its payment application on AWS with minimal operational overhead and infrastructure management. By using Amazon API Gateway, the company can create a secure and scalable API to receive payment notifications from mobile devices. By using AWS Lambda, the company can run a serverless function to validate the payment notifications and send them to the backend application. Lambda handles the provisioning, scaling, and security of the function, reducing the operational complexity and cost. By using Amazon ECS with AWS Fargate, the company can run the backend application on a fully managed container service that scales the compute resources automatically and does not require any EC2 instances to manage. Fargate allocates the right amount of CPU and memory for each container and adjusts them as needed.
* A. Create an Amazon Simple Queue Service (Amazon SQS) queue Integrate the queue with an Amazon EventBndge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS) Anywhere Create a standalone cluster. This option is not optimal because it requires the company to manage the Kubernetes cluster that runs the backend application. Amazon EKS Anywhere is a deployment option that allows the company to create and operate Kubernetes clusters on-premises or in other environments outside AWS. The company would need to provision, configure, scale, patch, and monitor the cluster nodes, which can increase the operational overhead and complexity. Moreover, the company would need to ensure the connectivity and security between the AWS services and the EKS Anywhere cluster, which can also add challenges and risks.
* B. Create an Amazon API Gateway API Integrate the API with anAWS Step Functions state ma-chine to receive payment notifications from mobile devices Invoke the state machine to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Sen/ice (Amazon EKS). Configure an EKS cluster with self-managed nodes. This option is not ideal because it requires the company to manage the EC2 instances that host the Kubernetes cluster that runs the backend application. Amazon EKS is a fully managed service that runs Kubernetes on AWS, but it still requires the company to manage the worker nodes that run the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, using AWS Step Functions to validate the payment notifications may be unnecessary and complex, as the validation logic can be implemented in a simpler way with Lambda or other services.
* C. Create an Amazon Simple Queue Sen/ice (Amazon SQS) queue Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon EC2 Spot Instances Configure a Spot Fleet with a default al-location strategy. This option is not cost-effective because it requires the company to manage the EC2 instances that run the backend application. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, using Spot Instances can introduce the risk of interruptions, as Spot Instances are reclaimed by AWS when the demand for On-Demand Instances increases. The company would need to handle the interruptions gracefully and ensure the availability and reliability of the backend application.
References:
? 1 Amazon API Gateway - Amazon Web Services
? 2 AWS Lambda - Amazon Web Services
? 3 Amazon Elastic Container Service - Amazon Web Services
? 4 AWS Fargate - Amazon Web Services

NEW QUESTION 105
- (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 109
- (Topic 4)
A company runs a three-tier web application in a VPC across multiple Availability Zones. Amazon EC2 instances run in an Auto Scaling group for the application tier.
The company needs to make an automated scaling plan that will analyze each resource's daily and weekly historical workload trends. The configuration must scale resources appropriately according to both the forecast and live changes in utilization.
Which scaling strategy should a solutions architect recommend to meet these requirements?

A. Implement dynamic scaling with step scaling based on average CPU utilization from the EC2 instances.
B. Enable predictive scaling to forecast and scal
C. Configure dynamic scaling with target tracking.
D. Create an automated scheduled scaling action based on the traffic patterns of the web application.
E. Set up a simple scaling polic
F. Increase the cooldown period based on the EC2 instance startup time

Answer: B

Explanation:
This solution meets the requirements because it allows the company to use both predictive scaling and dynamic scaling to optimize the capacity of its Auto Scaling group. Predictive scaling uses machine learning to analyze historical data and forecast future traffic patterns. It then adjusts the desired capacity of the group in advance of the predicted changes. Dynamic scaling uses target tracking to maintain a specified metric (such as CPU utilization) at a target value. It scales the group in or out as needed to keep the metric close to the target. By using both scaling methods, the company can benefit from faster, simpler, and more accurate scaling that responds to both forecasted and live changes in utilization. References:
? Predictive scaling for Amazon EC2 Auto Scaling
? [Target tracking scaling policies for Amazon EC2 Auto Scaling]

NEW QUESTION 114
- (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 117
- (Topic 4)
A medical research lab produces data that is related to a new study. The lab wants to make the data available with minimum latency to clinics across the country for their on-premises, file-based applications. The data files are stored in an Amazon S3 bucket that has read- only permissions for each clinic.
What should a solutions architect recommend to meet these requirements?

A. Deploy an AWS Storage Gateway file gateway as a virtual machine (VM) on premises at each clinic
B. Migrate the files to each clinic's on-premises applications by using AWS DataSync for processing.
C. Deploy an AWS Storage Gateway volume gateway as a virtual machine (VM) on premises at each clinic.
D. Attach an Amazon Elastic File System (Amazon EFS) file system to each clinic's on- premises servers.

Answer: A

Explanation:
AWS Storage Gateway is a service that connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure. By deploying a file gateway as a virtual machine on each clinic's premises, the medical research lab can provide low-latency access to the data stored in the S3 bucket while maintaining read-only permissions for each clinic. This solution allows the clinics to access the data files directly from their on-premises file-based applications without the need for data transfer or migration.

NEW QUESTION 122
- (Topic 4)
A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should.be protected throughout the entire application stack, and access to the information should be restricted to certain applications.
Which action should the solutions architect take?

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

Answer: C

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

NEW QUESTION 123
- (Topic 4)
A company wants to experiment with individual AWS accounts for its engineer team. The company wants to be notified as soon as the Amazon EC2 instance usage for a given month exceeds a specific threshold for each account.
What should a solutions architect do to meet this requirement MOST cost-effectively?

A. Use Cost Explorer to create a daily report of costs by servic
B. Filter the report by EC2instance
C. Configure Cost Explorer to send an Amazon Simple Email Service (Amazon SES) notification when a threshold is exceeded.
D. Use Cost Explorer to create a monthly report of costs by servic
E. Filter the report by EC2 instance
F. Configure Cost Explorer to send an Amazon Simple Email Service (Amazon SES) notification when a threshold is exceeded.
G. Use AWS Budgets to create a cost budget for each accoun
H. Set the period to monthl
I. Set the scope to EC2 instance
J. Set an alert threshold for the budge
K. Configure an Amazon Simple Notification Service (Amazon SNS) topic to receive a notification when a threshold is exceeded.
L. Use AWS Cost and Usage Reports to create a report with hourly granularit
M. Integrate the report data with Amazon Athen
N. Use Amazon EventBridge to schedule an Athena quer
O. Configure an Amazon Simple Notification Service (Amazon SNS) topic to receive a notification when a threshold is exceeded.

Answer: C

Explanation:
AWS Budgets allows you to create budgets for your AWS accounts and set alerts when usage exceeds a certain threshold. By creating a budget for each account, specifying the period as monthly and the scope as EC2 instances, you can effectively track the EC2 usage for each account and be notified when a threshold is exceeded. This solution is the most cost-effective option as it does not require additional resources such as Amazon Athena or Amazon EventBridge.

NEW QUESTION 128
- (Topic 4)
A company hosts multiple applications on AWS for different product lines. The applications use different compute resources, including Amazon EC2 instances and Application Load Balancers. The applications run in different AWS accounts under the same organization in
AWS Organizations across multiple AWS Regions. Teams for each product line have tagged each compute resource in the individual accounts.
The company wants more details about the cost for each product line from the consolidated billing feature in Organizations.
Which combination of steps will meet these requirements? (Select TWO.)

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

Answer: BE

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

NEW QUESTION 131
- (Topic 4)
A development team is collaborating with another company to create an integrated product. The other company needs to access an Amazon Simple Queue Service (Amazon SQS) queue that is contained in the development team's account. The other company wants to poll the queue without giving up its own account permissions to do so.
How should a solutions architect provide access to the SQS queue?

A. Create an instance profile that provides the other company access to the SQS queue.
B. Create an 1AM policy that provides the other company access to the SQS queue.
C. Create an SQS access policy that provides the other company access to the SQS queue.
D. Create an Amazon Simple Notification Service (Amazon SNS) access policy that provides the other company access to the SQS queue.

Answer: C

Explanation:
To provide access to the SQS queue to the other company without giving up its own account permissions, a solutions architect should create an SQS access policy that provides the other company access to the SQS queue. An SQS access policy is a resource-based policy that defines who can access the queue and what actions they can perform. The policy can specify the AWS account ID of the other company as a principal, and grant permissions for actions such as sqs:ReceiveMessage, sqs:DeleteMessage, and sqs:GetQueueAttributes. This way, the other company can poll the queue using its own credentials, without needing to assume a role or use cross-account access
keys. References:
? Using identity-based policies (IAM policies) for Amazon SQS
? Using custom policies with the Amazon SQS access policy language

NEW QUESTION 133
- (Topic 4)
A company has a financial application that produces reports. The reports average 50 KB in size and are stored in Amazon S3. The reports are frequently accessed during the first week after production and must be stored for several years. The reports must be retrievable within 6 hours.
Which solution meets these requirements MOST cost-effectively?

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

Answer: A

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

NEW QUESTION 134
- (Topic 4)
A company wants to migrate its on-premises Microsoft SQL Server Enterprise edition database to AWS. The company's online application uses the database to process transactions. The data analysis team uses the same production database to run reports for analytical processing. The company wants to reduce operational overhead by moving to managed services wherever possible.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: A

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

NEW QUESTION 135
- (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 138
- (Topic 4)
A company hosts an internal serverless application on AWS by using Amazon API Gateway and AWS Lambda. The company's employees report issues with high latency when they begin using the application each day. The company wants to reduce latency.
Which solution will meet these requirements?

A. Increase the API Gateway throttling limit.
B. Set up a scheduled scaling to increase Lambda provisioned concurrency before employees begin to use the application each day.
C. Create an Amazon CloudWatch alarm to initiate a Lambda function as a target for the alarm at the beginning of each day.
D. Increase the Lambda function memory.

Answer: B

Explanation:
AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda scales automatically based on the incoming requests, but it may take some time to initialize new instances of your function if there is a sudden increase in demand. This may result in high latency or cold starts for your application. To avoid this, you can use provisioned concurrency, which ensures that your function is initialized and ready to respond at any time. You can also set up a scheduled scaling policy that increases the provisioned concurrency before employees begin to use the application each day, and decreases it when the demand is low. References: https://docs.aws.amazon.com/lambda/latest/dg/configuration-concurrency.html

NEW QUESTION 143
- (Topic 4)
A company stores text files in Amazon S3. The text files include customer chat messages, date and time information, and customer personally identifiable information (Pll).
The company needs a solution to provide samples of the conversations to an external service provider for quality control. The external service provider needs to randomly pick sample conversations up to the most recent conversation. The company must not share the customer Pll with the external service provider. The solution must scale when the number of customer conversations increases.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Object Lambda Access Poin
B. Create an AWS Lambda function that redacts the Pll when the function reads the fil
C. Instruct the external service provider to access the Object Lambda Access Point.
D. Create a batch process on an Amazon EC2 instance that regularly reads all new files, redacts the Pll from the files, and writes the redacted files to a different S3 bucke
E. Instruct the external service provider to access the bucket that does not contain the Pll.
F. Create a web application on an Amazon EC2 instance that presents a list of the files, redacts the Pll from the files, and allows the external service provider to download new versions of the files that have the Pll redacted.
G. Create an Amazon DynamoDB tabl
H. Create an AWS Lambda function that reads only the data in the files that does not contain Pl
I. Configure the Lambda function to store the non-PII data in the DynamoDB table when a new file is written to Amazon S3. Grant the external service provider access to the DynamoDB table.

Answer: A

Explanation:
The correct solution is to create an Object Lambda Access Point and an AWS Lambda function that redacts the PII when the function reads the file. This way, the company can use the S3 Object Lambda feature to modify the S3 object content on the fly, without creating a copy or changing the original object. The external service provider can access the Object Lambda Access Point and get the redacted version of the file. This solution has the least operational overhead because it does not require any additional storage, processing, or synchronization. The solution also scales automatically with the number of customer conversations and the demand from the external service provider. The other options are incorrect because:
? Option B is using a batch process on an EC2 instance to read, redact, and write the files to a different S3 bucket. This solution has more operational overhead because it requires managing the EC2 instance, the batch process, and the additional S3 bucket. It also introduces latency and inconsistency between the original and the redacted files.
? Option C is using a web application on an EC2 instance to present, redact, and download the files. This solution has more operational overhead because it requires managing the EC2 instance, the web application, and the download process. It also exposes the original files to the web application, which increases the risk of leaking the PII.
? Option D is using a DynamoDB table and a Lambda function to store the non-PII
data from the files. This solution has more operational overhead because it requires managing the DynamoDB table, the Lambda function, and the data transformation. It also changes the format and the structure of the original files, which may affect the quality control process.
References:
? S3 Object Lambda
? Object Lambda Access Point
? Lambda function

NEW QUESTION 148
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Solution-Architect-Associate Exam with Our Prep Materials Via below:

https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html