iPassLeader

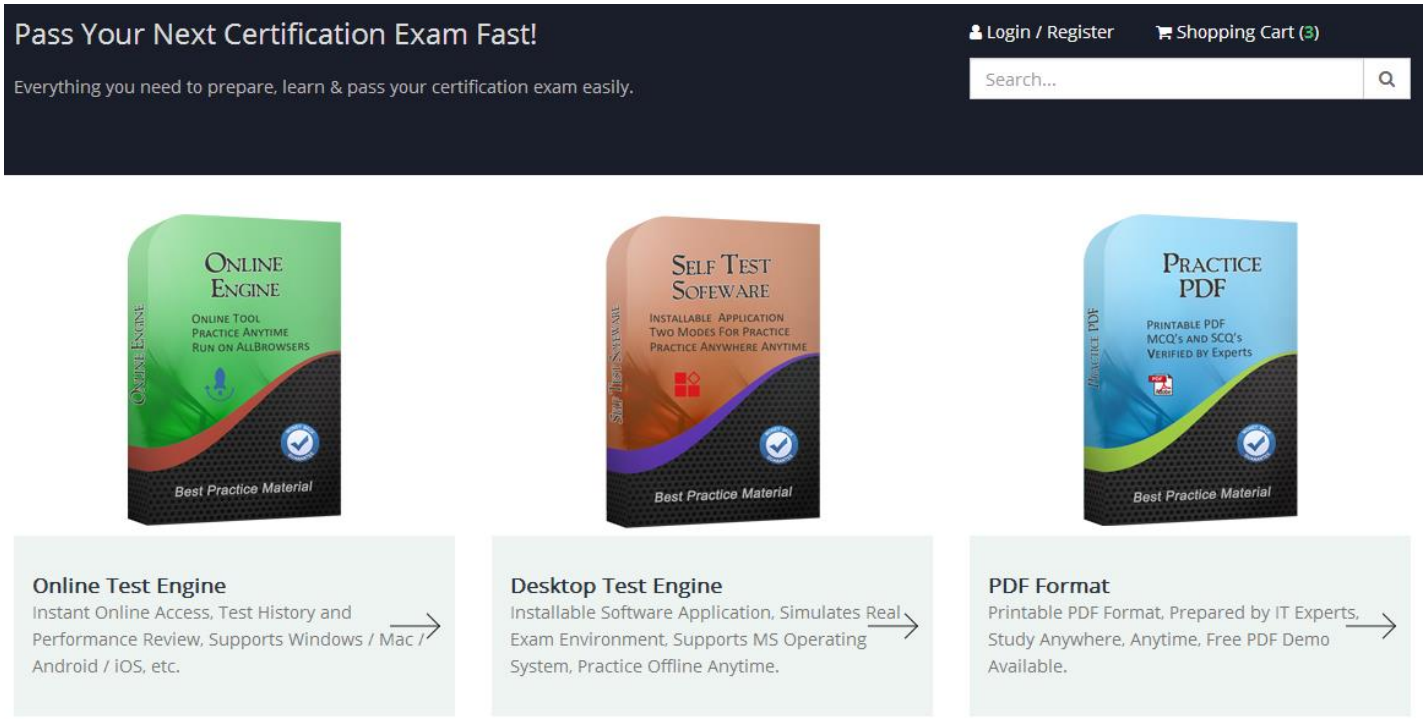

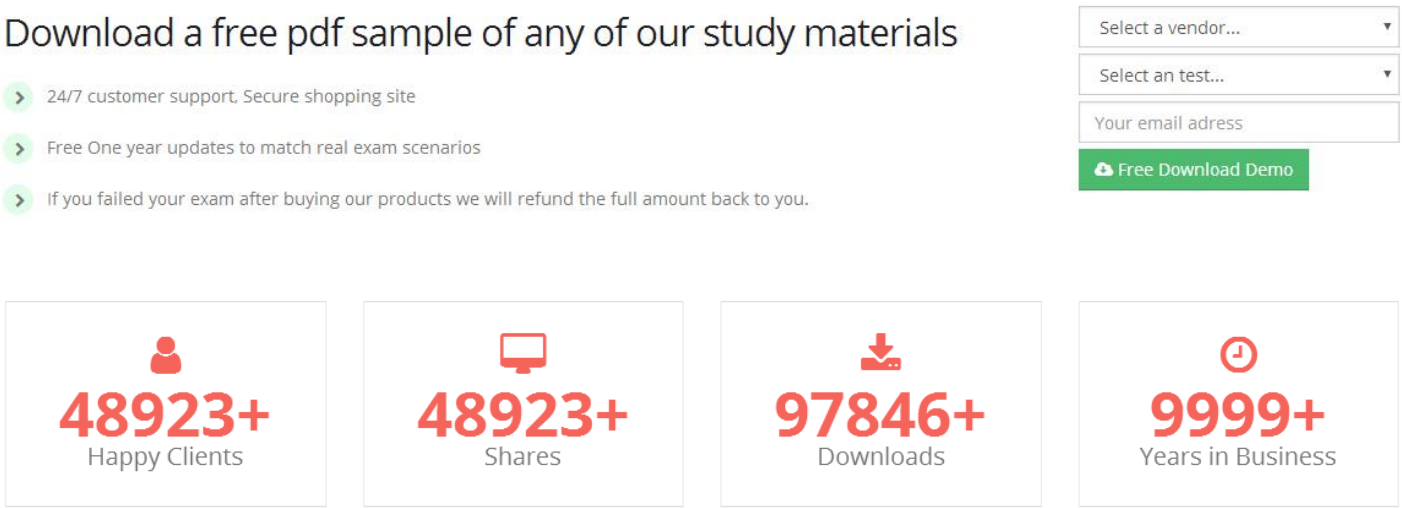

http://www.ipassleader.com/

Everything you need to prepare, learn & pass your certification exam easily.

Exam : **SAP-C02**

Title : AWS Certified Solutions Architect - Professional (SAP-C02)

Vendor : Amazon

Version : DEMO

QUESTION NO: 1
A company is designing a new website that hosts static content. The website will give users the ability to upload and download large files. According to company requirements, all data must be encrypted in transit and at rest. A solutions architect is building the solution by using Amazon S3 and Amazon CloudFront.
Which combination of steps will meet the encryption requirements? (Select THREE.)
A. Turn on S3 server-side encryption for the S3 bucket that the web application uses.
B. Add a policy attribute of " aws:SecureTransport " : " true " for read and write operations in the S3 ACLs.
C. Create a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses.
D. Configure encryption at rest on CloudFront by using server-side encryption with AWS KMS keys (SSE- KMS).
E. Configure redirection of HTTP requests to HTTPS requests in CloudFront.
F. Use the RequireSSL option in the creation of presigned URLs for the S3 bucket that the web application uses.
Answer: A C E
Explanation:
Turning on S3 server-side encryption for the S3 bucket that the web application uses will enable encrypting the data at rest using Amazon S3 managed keys (SSE-S3)1. Creating a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses will enable enforcing encryption for all requests to the bucket2. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS3.

QUESTION NO: 2
A utility company collects usage data from smart meters every 5 minutes. Data is sent to API Gateway, processed by Lambda, and stored in DynamoDB. As usage increased, Lambda durations increased and DynamoDB PUTs failed with ProvisionedThroughputExceededException. Lambda also experiences TooManyRequestsException errors.
Which combination of changes will resolve these issues? (Select TWO.)
A. Increase the payload size from the smart meters.
B. Collect data in an Amazon SQS FIFO queue, which triggers a Lambda function to process each message.
C. Stream the data into an Amazon Kinesis data stream from API Gateway and process the data in batches.
D. Increase the write capacity units to the DynamoDB table.
E. Increase the memory available to the Lambda functions.
Answer: C,D

QUESTION NO: 3
An application is using an Amazon RDS for MySQL Multi-AZ DB instance in the us-east-1 Region. After a failover test, the application lost the connections to the database and could

not re-establish the connections.
After a restart of the application, the application re-established the connections.
A solutions architect must implement a solution so that the application can re-establish connections to the database without requiring a restart.
Which solution will meet these requirements?

A. Create a two-node Amazon Aurora MySQL DB cluster. Migrate the RDS DB instance to the Aurora DB cluster. Create an RDS proxy. Configure the existing RDS endpoint as a target. Update the connection settings in the application to point to the RDS proxy endpoint.

B. Create an Amazon S3 bucket. Export the database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Configure Amazon Athena to use the S3 bucket as a data store. Install the latest Open Database Connectivity (ODBC) driver for the application. Update the connection settings in the application to point to the Athena endpoint

C. Create an Amazon Aurora MySQL Serverless v1 DB instance. Migrate the RDS DB instance to the Aurora Serverless v1 DB instance. Update the connection settings in the application to point to the Aurora reader endpoint.

D. Create an RDS proxy. Configure the existing RDS endpoint as a target. Update the connection settings in the application to point to the RDS proxy endpoint.

Answer: D

QUESTION NO: 4

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

A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.

B. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enabled. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template. Use the EBS volume as the shared storage for the duration of the job. Detach the EBS volume when the job is complete.

C. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses

the S3 Standard storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.

D. Migrate the data from the existing shared file system to an Amazon S3 bucket. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the job. Delete the file gateway when the job is complete.

Answer: A

Explanation:

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

QUESTION NO: 5

A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling Group The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted The maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:

VPCCIDR 10 0 0 0/23

AZ1 subnet CIDR: 10 0 0 0724

AZ2 subnet CIDR: 10.0.1 0724

Since deployment, a third AZ has become available in the Region The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime. Which solution will meet these requirements?

A. Update the Auto Scaling group to use the AZ2 subnet only Delete and re-create the AZ1 subnet using half the previous address space Adjust the Auto Scaling group to also use the new AZI subnet When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only Remove the current AZ2 subnet Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

B. Terminate the EC2 instances in the AZ1 subnet Delete and re-create the AZ1 subnet using hall the address space. Update the Auto Scaling group to use this new subnet. Repeat this for the second AZ.

Define a new subnet in AZ3: then update the Auto Scaling group to target all three new subnets

C. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ Update the existing Auto Scaling group to target the new subnets in the new VPC

D. Update the Auto Scaling group to use the AZ2 subnet only Update the AZ1 subnet to have halt the previous address space Adjust the Auto Scaling group to also use the AZ1 subnet again. When the instances are healthy, adjust the Auto Seating group to use the AZ1 subnet only. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet Create a new AZ3 subnet usinghalf the original AZ2 subnet address

space, then update the Auto Scaling group to target all three new subnets

Answer: A

Explanation:

https://repost.aws/knowledge-center/vpc-ip-address-range

QUESTION NO: 6

A healthcare company is building a user support chat-based assistant on Amazon Bedrock. Users will ask health questions that might include personal details in the prompts.

A solutions architect must configure a solution that can do the following:

* Prevent the assistant from providing medical diagnosis advice.
* Redact personally identifiable information (PII) from both user inputs and model responses.
* Enforce the same controls even if the company changes foundation models (FMs) later.
* Evaluate risky user prompts before sending the prompts to a model to avoid unnecessary inference costs.

Which solution will meet these requirements?

A. Store approved health support guidelines in an Amazon Bedrock knowledge base. Configure system prompts that instruct the model not to provide diagnosis advice. Use an AWS Lambda function after inference to remove PII from the model response before the response is returned to users.

B. Fine-tune an FM on approved support conversations. Add prompt templates that prohibit diagnosis advice. Run a separate review process that scans transcripts for prohibited topics and sensitive information after the conversations end.

C. Build a custom moderation layer in the application to inspect prompts for prohibited topics and to redact sensitive information from user inputs. Invoke the model through the Converse API. Use separate post-processing logic to redact sensitive information from responses before displaying responses to users.

D. Create an Amazon Bedrock guardrail. Configure denied topics for medical diagnosis advice. Configure sensitive information filters to mask PII. Configure content filters. Call the ApplyGuardrail API on user prompts before inference. Include the same guardrail in the Converse API to evaluate model responses.

Answer: D

Explanation:

Amazon Bedrock Guardrails are the correct managed control plane for this scenario. A guardrail can combine denied topics, content filters, and sensitive information filters, and it can be applied to prompts and responses across supported foundation models. The ApplyGuardrail API can evaluate text independently before invoking a foundation model, which helps reject or mask risky prompts before incurring model inference cost.

The Converse API also supports guardrail configuration so the same policy can evaluate conversational model responses. Option A relies on prompts and post-processing only, so it does not evaluate risky prompts before inference and does not protect inputs. Option B is retrospective and model-specific. Option C could work technically but creates custom moderation and response-filtering code, which is higher operational overhead than Bedrock Guardrails.

QUESTION NO: 7

A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.
How can this be accomplished?

A. Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda function. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.

B. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify code. Rollback if Amazon CloudWatch alarms are triggered.

C. Refactor the AWS CLI scripts into a single script that deploys the new Lambda version. When deployment is completed, the script tests execute. If errors are detected, revert to the previous Lambda version.

D. Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda version. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

Answer: B

Explanation:
https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and-phased- deployments-with-aws-codedeploy/

QUESTION NO: 8

A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration The company needs to store large, important documents within the application with the following requirements
1 The data must be highly durable and available
2. The data must always be encrypted at rest and in transit.
3 The encryption key must be managed by the company and rotated periodically Which of the following solutions should the solutions architect recommend?

A. Deploy the storage gateway to AWS in file gateway mode Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes

B. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.

C. Use Amazon DynamoDB with SSL to connect to DynamoDB Use an AWS KMS key to encrypt DynamoDB objects at rest.

D. Deploy instances with Amazon EBS volumes attached to store this data Use EBS volume encryption using an AWS KMS key to encrypt the data.

Answer: B

QUESTION NO: 9

A company is hosting a critical application on a single Amazon EC2 instance. The application uses an Amazon ElastiCache for Redis single-node cluster for an in-memory data store. The application uses an Amazon RDS for MariaDB DB instance for a relational database. For the application to function, each piece of the infrastructure must be healthy and must be in an active state.

A solutions architect needs to improve the application ' s architecture so that the infrastructure can automatically recover from failure with the least possible downtime.

Which combination of steps will meet these requirements? (Select THREE.)

A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.

B. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances Ensure that the EC2 instances are configured in unlimited mode.

C. Modify the DB instance to create a read replica in the same Availability Zone. Promote the read replica to be the primary DB instance in failure scenarios.

D. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.

E. Create a replication group for the ElastiCache for Redis cluster. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.

F. Create a replication group for the ElastiCache for Redis cluster. Enable Multi-AZ on the cluster.

Answer: A D F

Explanation:

Option A is correct because using an Elastic Load Balancer and an Auto Scaling group with a minimum capacity of two instances can improve the availability and scalability of the EC2 instances that host the application. The load balancer can distribute traffic across multiple instances and the Auto Scaling group can replace any unhealthy instances automatically1 Option D is correct because modifying the DB instance to create a Multi-AZ deployment that extends across two Availability Zones can improve the availability and durability of the RDS for MariaDB database. Multi- AZ deployments provide enhanceddata protection and minimize downtime by automatically failing over to astandby replica in another Availability Zone in case of a planned or unplanned outage4 Option F is correct because creating a replication group for the ElastiCache for Redis cluster and enabling Multi-AZ on the cluster can improve the availability and fault tolerance of the in-memory data store. A replication group consists of a primary node and up to five read-only replica nodes that are synchronized with the primary node using asynchronous replication. Multi-AZ allows automatic failover to one of the replicas if the primary node fails or becomes unreachable6 References: 1: https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/how-elastic-load-balancing- works.html 2: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/burstable-performance-instances- unlimited-mode.html 3: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html 4: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html

5: https://docs.
aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoScaling.html 6:
https://docs.aws.amazon.com
/AmazonElastiCache/latest/red-ug/Replication.Redis.Groups.html

QUESTION NO: 10
A company wants to establish a dedicated connection between its on-premises infrastructure and AWS. The company is setting up a 1 Gbps AWS Direct Connect connection to its account VPC. The architecture includes a transit gateway and a Direct Connect gateway to connect multiple VPCs and the on-premises infrastructure.
The company must connect to VPC resources over a transit VIF by using the Direct Connect connection.
Which combination of steps will meet these requirements? (Select TWO.)
A. Update the 1 Gbps Direct Connect connection to 10 Gbps.
B. Advertise the on-premises network prefixes over the transit VIF.
C. Adverse the VPC prefixes from the Direct Connect gateway to the on-premises network over the transit VIF.
D. Update the Direct Connect connection ' s MACsec encryption mode attribute to must encrypt.
E. Associate a MACsec Connection Key Name-Connectivity Association Key (CKN/CAK) pair with the Direct Connect connection.
Answer: B C
Explanation:
To connect VPC resources over a transit Virtual Interface (VIF) using a Direct Connect connection, the company should advertise the on-premises network prefixes over the transit VIF and advertise the VPC prefixes from the Direct Connect gateway to the on-premises network over the same VIF. This configuration ensures seamless connectivity between the on-premises infrastructure and the AWS VPCs through the transit gateway, facilitating efficient and secure communication across the network.
AWS Documentation on AWS Direct Connect and transit gateways provides detailed instructions on configuring transit VIFs and routing for Direct Connect connections. This setup is recommended in AWS best practices for establishing dedicated network connections between on-premises environments and AWS to achieve low-latency, high-throughput, and secure connectivity.

QUESTION NO: 11
A company plans to deploy a new private intranet service on Amazon EC2 instances inside a VPC. An AWS Site-to-Site VPN connects the VPC to the company ' s on-premises network. The new service must communicate with existing on-premises services The on-premises services are accessible through the use of hostnames that reside in the company example DNS zone This DNS zone is wholly hosted on premises and is available only on the company ' s private network.
A solutions architect must ensure that the new service can resolve hostnames on the company example domain to integrate with existing services.
Which solution meets these requirements?

A. Create an empty private zone in Amazon Route 53 for company example Add an additional NS record to the company ' s on-premises company example zone that points to the authoritative name servers for the new private zone in Route 53
B. Turn on DNS hostnames for the VPC Configure a new outbound endpoint with Amazon Route 53 Resolver. Create a Resolver rule to forward requests for company example to the on-premises name servers
C. Turn on DNS hostnames for the VPC Configure a new inbound resolver endpointwith Amazon Route
53 Resolver. Configure the on-premises DNS server to forward requests for company example to the new resolver.
D. Use AWS Systems Manager to configure a run document that will install a hosts file that contains any required hostnames. Use an Amazon EventBndge rule to run the document when an instance is entering the running state.
Answer: B
Explanation:
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver.html

QUESTION NO: 12
A company runs a workload in the AWS Cloud. The company stores data for the application in an older version of Amazon DocumentDB. Several backend services read and write data to the database continuously throughout all hours of the day. All services connect to the database by using the Amazon DocumentDB cluster endpoint, which is registered as a DNS record in Amazon Route 53.
The company needs to upgrade the database to the latest version of Amazon DocumentDB without losing any data. The company must be able to test and verify the upgrade before the company allows backend services to use the upgraded version. The company has already enabled change streams and set a retention period of 24 hours.
Which solution will meet these requirements?
A. Create a new Amazon DocumentDB cluster that runs the latest version. Use the Amazon DocumentDB Index Tool to export existing indexes and import them to the new cluster. Create a new AWS DMS instance and a source and target endpoint. Create a DMS task to migrate the data by using the Migrate and replicate migration type. Test and verify the new cluster. Update the Route 53 record to point to the new cluster.
B. Create a new Amazon DocumentDB cluster that runs the latest version. Install MongoDB command line interface (CLI) database tools on an Amazon EC2 instance. Use the MongoDB CLI to create a binary export, and import the data to the new Amazon DocumentDB cluster. Test and verify the new cluster. Update the Route 53 record to point to the new cluster.
C. Create a snapshot of the existing Amazon DocumentDB cluster. Perform an in-place major version upgrade. Modify the existing cluster to the latest version and the latest cluster parameter group. Apply modifications immediately. Test and verify the upgrade.
D. Create a new Amazon DocumentDB cluster that runs the latest version. Deploy the AWS DataSync agent to an Amazon EC2 instance and activate the agent. Create a new AWS DataSync task in enhanced mode. Start the transfer task to copy data to the new cluster. Test and verify the new cluster.Update the Route 53 record to point to the new cluster.

Answer: A
Explanation:
The company needs to upgrade DocumentDB to the latest version with no data loss while allowing continuous reads and writes. The company also must be able to test and verify the upgrade before switching production traffic. This is a classic requirement for performing an upgrade using a blue/green approach: build a new target environment on the new version, keep it in sync with the source, validate it, and then cut over by changing the endpoint (here, Route 53 DNS).
Option A implements this pattern using a new DocumentDB cluster running the latest version and AWS DMS to continuously migrate and replicate changes from the old cluster to the new cluster. Because the workload is continuously changing, a one-time export/import is insufficient; continuous replication is needed to keep the target cluster current during the test period. AWS DMS supports a "migrate and replicate" style of task that performs a full load and then applies ongoing changes (CDC) so the target stays synchronized. The question also states that change streams are enabled with a 24-hour retention period, which supports capturing and applying changes during migration/validation and helps ensure the replication stream can be maintained while testing.
Option A also addresses indexes by using the DocumentDB Index Tool to export and import indexes, which is important because indexes can affect query performance and behavior. After the company validates the new cluster, the cutover is done by updating the Route 53 record to point to the new cluster endpoint, switching all backend services without changing application configuration beyond DNS resolution.
Option B uses MongoDB CLI tools to export/import. This is not suitable for continuous write workloads because export/import is a point-in-time operation and would require downtime or risk data divergence during the test period. It also adds more operational overhead and does not provide continuous replication for the duration of validation.
Option C performs an in-place major version upgrade. That does not satisfy the requirement to test and verify the upgrade before backend services use the upgraded version because the upgrade happens directly on the production cluster. Even though a snapshot exists for rollback, production is still exposed to the upgrade immediately, which violates the requirement for pre-cutover verification.
Option D is incorrect because AWS DataSync transfers files between storage systems such as NFS/SMB and AWS storage services. It is not a database migration or replication service and cannot copy a DocumentDB database in a way that preserves database semantics and supports continuous replication.
Therefore, creating a new DocumentDB cluster, keeping it synchronized using AWS DMS (supported by change stream retention), validating it, and then cutting over via Route 53 DNS update (option A) meets all requirements.
References:
AWS documentation on blue/green style database upgrades by migrating to a new cluster and cutting over via DNS.
AWS documentation on AWS DMS full load plus ongoing replication (CDC) patterns for minimizing downtime and maintaining target synchronization during validation.
AWS documentation on Amazon DocumentDB change streams and retention considerations for capturing ongoing changes during migration windows.

QUESTION NO: 13
A company wants to migrate to AWS. The company is running thousands of VMs in a VMware ESXi environment. The company has no configuration management database and has little Knowledge about the utilization of the VMware portfolio.
A solutions architect must provide the company with an accurate inventory so that the company can plan for a cost-effective migration.
Which solution will meet these requirements with the LEAST operational overhead?
A. Use AWS Systems Manager Patch Manager to deploy Migration Evaluator to each VM. Review the collected data in Amazon QuickSight. Identify servers that have high utilization. Remove the servers that have high utilization from the migration list. Import the data to AWS Migration Hub.
B. Export the VMware portfolio to a csv file. Check the disk utilization for each server. Remove servers that have high utilization. Export the data to AWS Application Migration Service. Use AWS Server Migration Service (AWS SMS) to migrate the remaining servers.
C. Deploy the Migration Evaluator agentless collector to the ESXi hypervisor. Review the collected data in Migration Evaluator. Identify inactive servers. Remove the inactive servers from the migration list.
Import the data to AWS Migration Hub.
D. Deploy the AWS Application Migration Service Agent to each VM. When the data is collected, use Amazon Redshift to import and analyze the data. Use Amazon QuickSight for data visualization.
Answer: C
Explanation:
https://aws.amazon.com/migration-evaluator/features/

QUESTION NO: 14
A solutions architect is creating an AWS CloudFormation template from an existing manually created non- production AWS environment The CloudFormation template can be destroyed and recreated as needed The environment contains an Amazon EC2 instance The EC2 instance has an instance profile that the EC2 instance uses to assume a role in a parent account The solutions architect recreates the role in a CloudFormation template and uses the same role name When the CloudFormation template is launched in the child account, the EC2 instance can no longer assume the role in the parent account because of insufficient permissions What should the solutions architect do to resolve this issue?
A. In the parent account edit the trust policy for the role that the EC2 instance needs to assume Ensure that the target role ARN in the existing statement that allows the sts AssumeRole action is correct Save the trust policy
B. In the parent account edit the trust policy for the role that the EC2 instance needs to assume Add a statement that allows the sts AssumeRole action for the root principal of the child account Save the trust policy
C. Update the CloudFormation stack again Specify only the CAPABILITY_NAMED_IAM capability
D. Update the CloudFormation stack again Specify the CAPABIUTYJAM capability and the CAPABILITY_NAMEDJAM capability

Answer: A
Explanation:
Edit the Trust Policy:
Go to the IAM console in the parent account and locate the role that the EC2 instance needs to assume.
Edit the trust policy of the role to ensure that it correctly allows the sts action for the role ARN in the child account.
Update the Role ARN:
Verify that the target role ARN specified in the trust policy matches the role ARN created by the CloudFormation stack in the child account.
If necessary, update the ARN to reflect the correct role in the child account.
Save and Test:
Save the updated trust policy and ensure there are no syntax errors.
Test the setup by attempting to assume the role from the EC2 instance in the child account.
Verify that the instance can successfully assume the role and perform the required actions.
This ensures that the EC2 instance in the child account can assume the role in the parent account, resolving the permission issue.
References
AWS IAM Documentation on Trust Policies#51#.

QUESTION NO: 15
A company has loT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.
A solutions architect designs a solution in which the loT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically and are periodically encountering a RealProvisioned Throughput Exceeded error.
Which actions should the solution architect take to resolve this issue? (Select THREE.)
A. Reshard the stream to increase the number of shards s in the stream.
B. Use the Kinesis Producer Library KPL). Adjust the polling frequency.
C. Use consumers with the enhanced fan-out feature.
D. Reshard the stream to reduce the number of shards in the stream.
E. Use an error retry and exponential backoff mechanism in the consumer logic.
F. Configure the stream to use dynamic partitioning.
Answer: A C E
Explanation:
https://repost.aws/knowledge-center/kinesis-readprovisionedthroughputexceeded Follow Data Streams best practices To mitigate ReadProvisionedThroughputExceeded exceptions, apply these best practices:
* Reshard your stream to increase the number of shards in the stream.
* Use consumers with enhanced fan-out. For more information about enhanced fan-out, see Developing custom consumers with dedicated throughput (enhanced fan-out).
* Use an error retry and exponential backoff mechanism in the consumer logic if ReadProvisionedThroughputExceeded exceptions are encountered. For consumer applications that use an AWS SDK, the requests are retried by default.

QUESTION NO: 16
Question:
A company runs an application on Amazon EC2 and AWS Lambda. The application stores temporary data in Amazon S3. The S3 objects are deleted after 24 hours.
The company deploys new versions of the application by launching AWS CloudFormation stacks. The stacks create the required resources. After validating a new version, the company deletes the old stack. The deletion of an old development stack recently failed.
A solutions architect needs to resolve this issue without major architecture changes.
Which solution will meet these requirements?
A. Create a Lambda function to delete objects from the S3 bucket. Add the Lambda function as a custom resource in the CloudFormation stack with a DependsOn attribute that points to the S3 bucket resource.
B. Modify the CloudFormation stack to attach a DeletionPolicy attribute with a value of Delete to the S3 bucket.
C. Update the CloudFormation stack to add a DeletionPolicy attribute with a value of Snapshot for the S3 bucket resource.
D. Update the CloudFormation template to create an Amazon EFS file system to store temporary files instead of Amazon S3. Configure the Lambda functions to run in the same VPC as the EFS file system.
Answer: A
Explanation:
CloudFormation cannot delete non-empty S3 buckets. OptionAallows you to create acustom Lambda resourcethat deletes all objects in the S3 bucket before the stack deletes it. The DependsOn ensures the bucket deletion occurs only after the Lambda has completed.
B: Adding DeletionPolicy: Delete does not resolve the issue if the bucket still contains objects.
C: Snapshot doesn ' t apply to S3 and won ' t help here.
D: Changing to Amazon EFS would require architectural changes, which are not allowed per requirements.
Reference:https://aws.amazon.com/blogs/devops/safely-delete-s3-buckets-using-aws-cloudformation/https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute- deletionpolicy.html

QUESTION NO: 17
A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company ' s AWS accounts.
Which combination of steps will meet these requirements? (Choose two.)
A. In the centralized account, create an IAM role that has the Lambda service as a trusted entity. Add an inline policy to assume the roles of the other AWS accounts.
B. In the other AWS accounts, create an IAM role that has minimal permissions. Add the centralized account ' s Lambda IAM role as a trusted entity.

C. In the centralized account, create an IAM role that has roles of the other accounts as trusted entities.
Provide minimal permissions.
D. In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized account. Add the Lambda service as a trusted entity.
E. In the other AWS accounts, create an IAM role that has minimal permissions. Add the Lambda service as a trusted entity.
Answer: A B
Explanation:
https://medium.com/@it.melnichenko/invoke-a-lambda-across-multiple-aws-accounts-8c094b2e70be

QUESTION NO: 18
Question:
A company has an application that stores user-uploaded videos in an Amazon S3 bucket using S3 Standard storage. Users access videos frequently for the first 180 days, and rarely after that. Most videos are over 100 MB. Users often have poor internet connectivity, and the company uses multipart uploads.
A solutions architect needs tooptimize S3 storage costs.
Which combination of actions will meet these requirements? (Select TWO.)
A. Configure the S3 bucket to be a Requester Pays bucket.
B. Use S3 Transfer Acceleration to upload the videos.
C. Create a lifecycle rule to expireincomplete multipart uploadsafter 7 days.
D. Create a lifecycle rule to transition objects toS3 Glacier Instant Retrieval after 1 day.
E. Create a lifecycle rule to transition objects toS3 Standard-IA after 180 days.
Answer: C E
Explanation:
C: Multipart uploads can leave incomplete parts behind, which incur storage costs. Expiring them after 7 days minimizes waste and saves cost.
E: Since objects are infrequently accessed after 180 days, transitioning toS3 Standard-IAis cost-effective, especially for large files > 128 KB (your 100 MB+ files qualify).
Ais for shifting download cost, not reducing your S3 storage expenses.
Bhelps with upload speed butincreases cost.
Dis too aggressive; Glacier is not suited for access patterns within the first few days.
Reference:https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-configuration-examples.
htmlhttps://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html

QUESTION NO: 19
A company is planning to migrate an on-premises data center to AWS. The company currently hosts the data center on Linux-based VMware VMs. A solutions architect must collect information about network dependencies between the VMs. The information must be in the form of a diagram that details host IP addresses, hostnames, and network connection information.
Which solution will meet these requirements?

A. Use AWS Application Discovery Service. Select an AWS Migration Hub home AWS Region. Install the AWS Application Discovery Agent on the on-premises servers for data collection. Grant permissions to Application Discovery Service to use the Migration Hub network diagrams.
B. Use the AWS Application Discovery Service Agentless Collector for server data collection. Export the network diagrams from the AWS Migration Hub in .png format.
C. Install the AWS Application Migration Service agent on the on-premises servers for data collection.
Use AWS Migration Hub data in Workload Discovery on AWS to generate network diagrams.
D. Install the AWS Application Migration Service agent on the on-premises servers for data collection.Export data from AWS Migration Hub in .csv format into an Amazon CloudWatch dashboard to generate network diagrams.
Answer: A
Explanation:
To effectively gather information about network dependencies between VMs in an on-premises data center for migration to AWS, it ' s crucial to use tools that can capture detailed application and server dependencies. The AWS Application Discovery Service is designed for this purpose, particularly when migrating from environments like Linux-based VMware VMs. By installing the AWS Application Discovery Agent on the on- premises servers, the service can collect necessary data such as host IP addresses, hostnames, and network connection information. This data is crucial for creating a comprehensive network diagram that outlines the interactions and dependencies between various components of the on-premises infrastructure. The integration with AWS Migration Hub enhances this process by allowing the visualization of these dependencies in a network diagram format, aiding in the planning and execution of the migration process. This approach ensures a thorough understanding of the on-premises environment, which is essential for a successful migration to AWS.
References:
AWS Documentation on Application Discovery Service: This provides detailed guidance on how to use the Application Discovery Service, including the installation and configuration of the Discovery Agent.
AWS Migration Hub User Guide: Offers insights on how to integrate Application Discovery Service data with Migration Hub for comprehensive migration planning and tracking.
AWS Solutions Architect Professional Learning Path: Contains advanced topics and best practices for migrating complex on-premises environments to AWS, emphasizing the use of AWS services and tools for effective migration planning and execution.

QUESTION NO: 20
A company is hosting an application on AWS for a project that will run for the next 3 years. The application consists of 20 Amazon EC2 On-Demand Instances that are registered in a target group for a Network Load Balancer (NLB). The instances are spread across two Availability Zones. The application is stateless and runs
24 hours a day, 7 days a week.
The company receives reports from users who are experiencing slow responses from the application.

Performance metrics show that the instances are at 10% CPU utilization during normal application use.
However, the CPU utilization increases to 100% at busy times, which typically last for a few hours.
The company needs a new architecture to resolve the problem of slow responses from the application.
Which solution will meet these requirements MOST cost-effectively?

A. Create an Auto Scaling group. Attach the Auto Scaling group to the target group of the NLB. Set the minimum capacity to 20 and the desired capacity to 28. Purchase Reserved Instances for 20 instances.

B. Create a Spot Fleet that has a request type of request. Set the TotalTargetCapacity parameter to 20. Set the DefaultTargetCapacityType parameter to On-Demand. Specify the NLB when creating the Spot Fleet.

C. Create a Spot Fleet that has a request type of maintain. Set the TotalTargetCapacity parameter to 20. Set the DefaultTargetCapacityType parameter to Spot. Replace the NLB with an Application Load Balancer.

D. Create an Auto Scaling group. Attach the Auto Scaling group to the target group of the NLB. Set the minimum capacity to 4 and the maximum capacity to 28. Purchase Reserved Instances for four instances.

Answer: D

QUESTION NO: 21

A company has developed a hybrid solution between its data center and AWS. The company uses Amazon VPC and Amazon EC2 instances that send application logs to Amazon CloudWatch. The EC2 instances read data from multiple relational databases that are hosted on premises.
The company wants to monitor which EC2 instances are connected to the databases in near real time. The company already has a monitoring solution that uses Splunk on premises. A solutions architect needs to determine how to send networking traffic to Splunk.
How should the solutions architect meet these requirements?

A. Enable VPC flow logs and send them to CloudWatch. Create an AWS Lambda function to periodically export the CloudWatch logs to an Amazon S3 bucket by using the predefined export function. Generate ACCESS_KEY and SECRET_KEY AWS credentials. Configure Splunk to pull the logs from the S3 bucket by using those credentials.

B. Create an Amazon Data Firehose delivery stream with Splunk as the destination. Configure a pre- processing AWS Lambda function with a Firehose stream processor that extracts individual log events from records sent by CloudWatch Logs subscription filters. Enable VPC flow logs and send them to CloudWatch. Create a CloudWatch Logs subscription that sends log events to the Firehose delivery stream.

C. Ask the company to log every request that is made to the databases along with the EC2 instance IP address. Export the CloudWatch logs to an Amazon S3 bucket. Use Amazon Athena to query the logs grouped by database name. Export Athena results to another S3 bucket. Invoke an AWS Lambda function to automatically send any new file that is put in the S3 bucket to Splunk.

D. Send the CloudWatch logs to an Amazon Kinesis data stream with Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics). Configure a 1-minute sliding window to collect the events. Create a SQL query that uses the anomaly detection template to monitor any networking traffic anomalies in near real time. Send the result to an Amazon Data Firehose delivery stream with Splunk as the destination.

Answer: B

Explanation:

The company needs near-real-time visibility into which EC2 instances are connecting to on-premises databases. The correct telemetry source for network connection metadata at the VPC level is VPC Flow Logs.

VPC Flow Logs capture information about IP traffic going to and from network interfaces in a VPC, including source/destination IPs, ports, protocol, and accept/reject decisions. This data can be used to infer which EC2 instance IPs are connecting to database IPs.

The company already uses Splunk on premises, so the solution should deliver these logs to Splunk with minimal delay and operational overhead. Amazon Data Firehose provides a fully managed way to deliver streaming data to supported destinations, including Splunk, with buffering and retry handling. CloudWatch Logs subscription filters can stream log events in near real time from CloudWatch Logs to destinations such as Firehose.

Option B uses the standard pattern: enable VPC Flow Logs to CloudWatch Logs, then create a CloudWatch Logs subscription filter that streams the flow logs to a Firehose delivery stream configured with Splunk as the destination. Because CloudWatch Logs subscription deliveries can batch log events, using a Firehose preprocessing Lambda to extract individual log events is a common approach to format records in a way that Splunk ingests cleanly. This yields near-real-time delivery with low operational overhead.

Option A introduces delay because it exports CloudWatch logs periodically to S3 and requires Splunk to poll S3. It also requires long-lived access keys and periodic batch exports, which is not near real time.

Option C relies on application-level logging changes and batch analytics with Athena, which is not near real time and requires substantial changes and additional pipelines.

Option D is over-engineered for the stated requirement. Using Flink and anomaly detection focuses on anomalies rather than simply identifying connections, and it adds significant operational complexity compared to direct delivery of flow logs to Splunk via Firehose.

Therefore, streaming VPC Flow Logs from CloudWatch Logs to Splunk using a Firehose delivery stream and a subscription filter is the best approach.

References:AWS documentation on VPC Flow Logs and the metadata they provide for network connection visibility.AWS documentation on CloudWatch Logs subscription filters for near-real-time streaming of log events.AWS documentation on Amazon Data Firehose delivery to Splunk and optional Lambda transformations for record formatting.

QUESTION NO: 22

A company completed a successful Amazon Workspaces proof of concept. They now want to make Workspaceshighly available across two AWS Regions. Workspaces are deployed in the failover Region. A hosted zone is available in Amazon Route 53.

What should the solutions architect do?

A. Create a connection alias in the primary Region and in the failover Region. Associate each

with a directory in its Region. Create a Route 53 failover routing policy with Evaluate Target Health = Yes.

B. Create a connection alias in both Regions. Associate both with a directory in the primary Region. Use a Route 53 multivalue answer routing policy.

C. Create a connection alias in the primary Region. Associate with the directory in the primary Region.
Use Route 53 weighted routing.

D. Create a connection alias in the primary Region. Associate it with the directory in the failover Region.Use Route 53 failover routing with Evaluate Target Health = Yes.

Answer: A

Explanation:
A is correct because AWS recommends using oneconnection alias per Region, associated witheach directory.
Then, configure a Route 53 failover policy so that if the primary Region becomes unhealthy, users are directed to the failover Region automatically. "Evaluate Target Health" ensures automatic detection and failover.
References:
Amazon Workspaces Cross-Region Resilience
Route 53 Failover Routing

QUESTION NO: 23

A company has an organization in AWS Organizations. The company is using AWS Control Tower to deploy a landing zone for the organization. The company wants to implement governance and policy enforcement.
The company must implement a policy that will detect Amazon RDS DB instances that are not encrypted at rest in the company's production OU.
Which solution will meet this requirement?

A. Turn on mandatory guardrails in AWS Control Tower. Apply the mandatory guardrails to the production OU.

B. Enable the appropriate guardrail from the list of strongly recommended guardrails in AWS Control Tower. Apply the guardrail to the production OU.

C. Use AWS Config to create a new mandatory guardrail. Apply the rule to all accounts in the production OU.

D. Create a custom SCP in AWS Control Tower. Apply the SCP to the production OU.

Answer: B

Explanation:
AWS Control Tower provides a set of " strongly recommended guardrails " that can be enabled to implement governance and policy enforcement. One of these guardrails is " Encrypt Amazon RDS instances " which will detect RDS DB instances that are not encrypted at rest. By enabling this guardrail and applying it to the production OU, the company will be able to enforce encryption for RDS instances in the production environment.

QUESTION NO: 24

An adventure company has launched a new feature on its mobile app. Users can use the feature to upload their hiking and ratting photos and videos anytime. The photos and videos

are stored in Amazon S3 Standard storage in an S3 bucket and are served through Amazon CloudFront.
The company needs to optimize the cost of the storage. A solutions architect discovers that most of the uploaded photos and videos are accessed infrequently after 30 days. However, some of the uploaded photos and videos are accessed frequently after 30 days. The solutions architect needs to implement a solution that maintains millisecond retrieval availability of the photos and videos at the lowest possible cost.
Which solution will meet these requirements?

A. Configure S3 Intelligent-Tiering on the S3 bucket.

B. Configure an S3 Lifecycle policy to transition image objects and video objects from S3 Standard to S3 Glacier Deep Archive after 30 days.

C. Replace Amazon S3 with an Amazon Elastic File System (Amazon EFS) file system that is mounted on Amazon EC2 instances.

D. Add a Cache-Control: max-age header to the S3 image objects and S3 video objects. Set the header to
30 days.

Answer: A
Explanation:
Amazon S3 Intelligent-Tiering is a storage class that automatically moves objects between two access tiers based on changing access patterns. Objects that are accessed frequently are stored in the frequent access tier and objects that are accessed infrequently are stored in the infrequent access tier. This allows for cost optimization without requiring manual intervention. This makes it an ideal solution for the scenario described, as it can automatically move objects that are infrequently accessed after 30 days to a lower-cost storage tier while still maintaining millisecond retrieval availability.

QUESTION NO: 25
A company has developed a new release of a popular video game and wants to make it available for public download The new release package is approximately 5 GB in size. The company provides downloads for existing releases from a Linux-based publicly facing FTP site hosted in an on-premises data center The company expects the new release will be downloaded by users worldwide The company wants a solution that provides improved download performance and low transfer costs regardless of a user ' s location Which solutions will meet these requirements ' ?

A. Store the game files on Amazon EBS volumes mounted on Amazon EC2 instances within an Auto Scaling group Configure an FTP service on the EC2 instances Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to downloadthe package

B. Store the game files on Amazon EFS volumes that are attached to Amazon EC2 instances within an Auto Scaling group Configure an FTP service on each of the EC2 instances Use an Application Load Balancer in front of the Auto Scaling group Publish the game download URL for users to download the package

C. Configure Amazon Route 53 and an Amazon S3 bucket for website hosting Upload the game files to the S3 bucket Use Amazon CloudFront for the website Publish the game download URL for users to download the package

D. Configure Amazon Route 53 and an Amazon S3 bucket for website hosting Upload the game files to the S3 bucket Set Requester Pays for the S3 bucket Publish the game download URL for users to download the package

Answer: C

Explanation:

Create an S3 Bucket:

Navigate to Amazon S3 in the AWS Management Console and create a new S3 bucket to store the game files.

Enable static website hosting on this bucket.

Upload Game Files:

Upload the 5 GB game release package to the S3 bucket. Ensure that the files are publicly accessible if required for download.

Configure Amazon Route 53:

Set up a new domain or subdomain in Amazon Route 53 and point it to the S3 bucket. This allows users to access the game files using a custom URL.

Use Amazon CloudFront:

Create a CloudFront distribution with the S3 bucket as the origin. CloudFront is a content delivery network (CDN) that caches content at edge locations worldwide, improving download performance and reducing latency for users regardless of their location.

Publish the Download URL:

Use the CloudFront distribution URL as the download link for users to access the game files. CloudFront will handle the efficient distribution and caching of the content.

This solution leverages the scalability of Amazon S3 and the performance benefits of CloudFront to provide an optimal download experience for users globally while minimizing costs.

References

Amazon CloudFront Documentation

Amazon S3 Static Website Hosting

QUESTION NO: 26

A company that has multiple AWS accounts is using AWS Organizations. The company's AWS accounts host VPCs, Amazon EC2 instances, and containers.

The company's compliance team has deployed a security tool in each VPC where the company has deployments. The security tools run on EC2 instances and send information to the AWS account that is dedicated for the compliance team. The company has tagged all the compliance-related resources with a key of "costCenter" and a value or "compliance".

The company wants to identify the cost of the security tools that are running on the EC2 instances so that the company can charge the compliance team's AWS account. The cost calculation must be as accurate as possible.

What should a solutions architect do to meet these requirements?

A. In the management account of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account.

Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.

B. In the member accounts of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account.
Schedule a monthly AWS Lambda function to retrieve the reports and calculate the total cost for the costCenter tagged resources.
C. In the member accounts of the organization activate the costCenter user-defined tag. From the management account, schedule a monthly AWS Cost and Usage Report. Use the tag breakdown in the report to calculate the total cost for the costCenter tagged resources.
D. Create a custom report in the organization view in AWS Trusted Advisor. Configure the report to generate a monthly billing summary for the costCenter tagged resources in the compliance team's AWS account.
Answer: A
Explanation:
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.htmlhttps://docs.aws.amazon.com
/awsaccountbilling/latest/aboutv2/configurecostallocreport.html

QUESTION NO: 27
A company runs a latency-sensitive application that consumes messages from an Amazon Managed Streaming for Apache Kafka (Amazon MSK) cluster. The MSK cluster runs across three Availability Zones.
The current MSK cluster uses Standard brokers with two standard large instances in each Availability Zone.
The company wants to minimize latency between Apache Kafka clients that are deployed in the same Availability Zones as the brokers. The company wants to increase available bandwidth and to increase the scaling speed of the cluster. Clients currently use default settings. Some downtime is acceptable while the company implements a solution.
Which solution will meet these requirements?
A. Configure a predictive scaling policy and set the MSK cluster as the target. Set the target value to 80 and set the scheduling buffer size to 0. Configure a placement group for the Kafka clients and associate the MSK hosts with the placement group.
B. Configure Cruise Control on the MSK cluster and enable bandwidth control bandwidth and rebalancing.
Deploy an Amazon MSK Connect proxy layer that uses latency-based routing. Reconfigure the Kafka clients to use the proxy endpoint.
C. Replace the Standard brokers with Express brokers that use express large instances. Set the client.rack property for the Kafka clients to az_id.
D. Resize the brokers to standard xlarge instances. Create MSK PrivateLink endpoints in each Availability Zone. Reconfigure each Kafka client to use the endpoint that is in the same Availability Zone as the client.
Answer: C
Explanation:
The company wants three things: minimize client-to-broker latency within the same Availability Zone, increase available bandwidth, and increase the scaling speed of the MSK

cluster. The current brokers are Standard brokers (two per AZ). Clients use default settings, which means they are not explicitly configured for rack awareness or AZ affinity.
A common way to reduce latency in multi-AZ Kafka deployments is to enable rack awareness on clients and brokers so clients prefer brokers in the same "rack," which can map to an Availability Zone. In Kafka, the client.rack setting allows the client to include rack information so the broker can return metadata that helps the client select replicas that are closest, reducing cross-AZ traffic and improving latency.
To increase bandwidth and improve scaling speed, the most direct approach in the choices is to move from Standard brokers to Express brokers. Express brokers are designed to provide higher throughput and faster scaling characteristics compared to standard broker types.
Since the question explicitly calls out increasing available bandwidth and scaling speed, the broker type change is the key lever, and it can be combined with client.rack configuration to minimize cross-AZ latency.
Option C matches these requirements: it replaces Standard brokers with Express brokers (to improve throughput/bandwidth and scaling speed) and sets client.rack to the Availability Zone identifier (az_id) to improve locality and reduce latency between clients and brokers in the same AZ.
Option A is not appropriate because MSK does not use EC2 Auto Scaling predictive scaling in that manner, and Kafka clients/brokers are not "associated" with an EC2 placement group as a primary latency solution in MSK. Placement groups are for EC2 instance placement; MSK broker placement is managed by the service.
Option B introduces a proxy layer and MSK Connect in a way that increases complexity and does not directly guarantee lower latency or higher bandwidth. MSK Connect is for Kafka Connect workloads, not as a general- purpose low-latency routing proxy for Kafka clients. Cruise Control is used for partition rebalancing and cluster optimization, but it does not replace the benefits of higher-throughput broker types and client rack awareness for AZ locality.
Option D increases broker size and introduces PrivateLink endpoints. PrivateLink is about private connectivity from VPCs to services and does not inherently ensure AZ-local broker selection or reduce latency between clients and brokers in the same AZ. Also, resizing to xlarge increases capacity but does not address scaling speed and locality as directly as express brokers plus rack configuration.
Therefore, option C best meets all requirements.
References:AWS documentation on Amazon MSK broker types, including performance and scaling characteristics of Standard and Express brokers.Apache Kafka concepts and AWS guidance on rack awareness and using client.rack to reduce cross-AZ traffic and latency in multi-AZ Kafka deployments.

QUESTION NO: 28
A company processes environment data. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.
The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time.
Which solution will meet these requirements?
A. Use Amazon Kinesis Data Firehouse to send the data to Amazon Redshift.

B. Use Amazon Kinesis Data streams to send the data to Amazon DynamoDB.
C. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to send the data to Amazon Aurora.
D. Use Amazon Kinesis Data firehouse to send the data to Amazon Keyspaces (for Apache Cassandra).
Answer: B
Explanation:
Amazon Kinesis Data Streams is a service that enables real-time data ingestion and processing. Amazon DynamoDB is a NoSQL database that does not require fixed schemas for storage. By using Kinesis Data Streams and DynamoDB, the company can send the JSON data to a database that can handle schemaless data in real time. References:
https://docs.aws.amazon.com/streams/latest/dev/introduction.html
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html

QUESTION NO: 29
A company is using GitHub Actions to run a CI/CD pipeline that accesses resources on AWS. The company has an IAM user that uses a secret key in the pipeline to authenticate to AWS. An existing IAM role with an attached policy grants the required permissions to deploy resources.
The company ' s security team implements a new requirement that pipelines can no longer use long-lived secret keys. A solutions architect must replace the secret key with a short-lived solution.
Which solution will meet these requirements with the LEAST operational overhead?
A. Create an IAM SAML 2.0 identity provider (IdP) in IAM. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRole API call. Attach the existing IAM policy to the new IAM role. Update GitHub to use SAML authentication for the pipeline.
B. Create an IAM OpenID Connect (OIDC) identity provider (IdP) in IAM. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRoleWithWebIdentity API call from the GitHub OIDC IdP. Update GitHub to assume the role for the pipeline.
C. Create an Amazon Cognito identity pool. Configure the authentication provider to use GitHub. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRoleWithWebIdentity API call from the GitHub authentication provider. Configure the pipeline to use Cognito as its authentication provider.
D. Create a trust anchor to AWS Private CA. Generate a client certificate to use with AWS IAM Roles Anywhere. Create a new IAM role with the appropriate trust policy that allows the sts:AssumeRole API call. Attach the existing IAM policy to the new IAM role. Configure the pipeline to use the credential helper tool and to reference the client certificate public key to assume the new IAM role.
Answer: B
Explanation:
This explanation is based on AWS documentation and best practices but is paraphrased, not a literal extract.
The current CI/CD pipeline uses an IAM user with long-lived access keys stored in GitHub Actions. The new requirement is that pipelines must not use long-lived secret keys. Instead,

the solution should provide short- lived credentials with minimal operational overhead.
GitHub Actions natively supports integration with cloud providers using OpenID Connect (OIDC). With OIDC, GitHub acts as an identity provider that can issue OIDC tokens to workflows. On the AWS side, IAM supports configuring an OIDC identity provider and roles that can be assumed by principals presenting valid OIDC tokens through the sts:AssumeRoleWithWebIdentity API. This pattern enables short-lived, automatically rotated credentials for CI/CD jobs without storing long-lived secrets.
In the correct solution (option B), you configure an IAM OIDC identity provider for GitHub in the AWS account. You then create a new IAM role with a trust policy that allows the Github OIDC provider to call sts:
AssumeRoleWithWebIdentity, with conditions that restrict which repositories or workflows can assume the role. The existing IAM policy that grants deployment permissions is attached to that role. In GitHub Actions, you update the pipeline configuration to request an OIDC token and assume the IAM role at runtime. Each workflow run receives short-lived credentials without storing static keys, and AWS automatically handles the token verification and temporary credential issuance. This approach is the AWS-recommended pattern for integrating GitHub Actions with AWS without long-lived secrets and has low operational overhead once configured.
Option A uses SAML 2.0, which is typically used for enterprise single sign-on for users, not for GitHub Actions workflows. GitHub does not natively use SAML to obtain AWS credentials for CI/CD pipelines in the same streamlined way as OIDC, and implementing a SAML-based integration would add unnecessary complexity.
Option C introduces Amazon Cognito as an indirection layer. Although Cognito can federate with external identity providers, including social providers, using it as an intermediary to obtain temporary AWS credentials for a machine-to-machine CI/CD pipeline is not necessary when IAM OIDC federation with GitHub is directly supported. This adds additional configuration and operational overhead.
Option D uses IAM Roles Anywhere with client certificates from AWS Private CA. Roles Anywhere is designed for workloads running outside AWS that need to assume IAM roles using X.509 certificates instead of access keys. While technically possible, it requires managing private certificates, trust anchors, and a credential helper tool, which is more complex and operationally heavier than the direct OIDC integration specifically designed for GitHub Actions.
Therefore, configuring an IAM OIDC identity provider for GitHub and creating an IAM role to be assumed via sts:AssumeRoleWithWebIdentity (option B) meets the requirement to replace long-lived secret keys with short-lived credentials with the least operational overhead .
References:AWS documentation on configuring IAM OpenID Connect identity providers and roles for GitHub Actions integration.AWS security best practices recommending federation and temporary credentials over long-lived IAM user access keys for CI/CD pipelines.

QUESTION NO: 30
A company has purchased appliances from different vendors. The appliances all have loT sensors. The sensors send status information in the vendors ' proprietary formats to a legacy application that parses the information into JSON. The parsing is simple, but each vendor has a unique format. Once daily, the application parses all the JSON records and stores the

records in a relational database for analysis.
The company needs to design a new data analysis solution that can deliver faster and optimize costs.
Which solution will meet these requirements?

A. Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the files. Use Amazon Athena and Amazon OuickSight for analysis.

B. Migrate the application server to AWS Fargate, which will receive the information from IoT sensors and parse the information into a relational format. Save the parsed information to Amazon Redshift for analysis.

C. Create an AWS Transfer for SFTP server. Update the IoT sensor code to send the information as a .csv file through SFTP to the server. Use AWS Glue to catalog the files. Use Amazon Athena for analysis.

D. Use AWS Snowball Edge to collect data from the IoT sensors directly to perform local analysis.Periodically collect the data into Amazon Redshift to perform global analysis.

Answer: A

Explanation:
Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the files. Use Amazon Athena and Amazon QuickSight for analysis. This solution meets the requirement of faster analysis and cost optimization by using AWS IoT Core to collect data from the IoT sensors in real-time and then using AWS Glue and Amazon Athena for efficient data analysis.
This solution involves connecting the IoT sensors to the AWS IoT Core, setting a rule to invoke an AWS Lambda function to parse the information, and saving a .csv file to Amazon S3. AWS Glue can be used to catalog the files and Amazon Athena and Amazon QuickSight can be used for analysis. This solution will enable faster and more cost-effective data analysis.
This solution is in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that: "AWS IoT Core can be used to ingest and process the data, AWS Lambda can be used to process and transform the data, and Amazon S3 can be used to store the data. AWS Glue can be used to catalog and access the data, Amazon Athena can be used to query the data, and Amazon QuickSight can be used to visualize the data." (Source:https://d1.awsstatic.com /training-and-certification/docs-sa-pro /AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf)

QUESTION NO: 31

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

Which solution is the MOST cost-effective way to meet these requirements?

A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in each account to create monthly reports for each business unit.

B. Configure AWS Budgets in the organization ' s master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization ' s master account to create monthly reports for each business unit.

C. Configure AWS Budgets in each account and configure budget alerts lhat are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.

D. Enable AWS Cost and Usage Reports in the organization ' s master account and configure reports grouped by application, environment, and owner. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit ' s email list.

Answer: B

Explanation:

Configure AWS Budgets in the organization # €™s master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization # €™s master account to create monthly reports for each business unit.

https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%
20Budgets%20gives%20you%20the,below%20the%20threshold%20you%20define

QUESTION NO: 32

A company is deploying a new web-based application and needs a storage solution for the Linux application servers. The company wants to create a single location for updates to application data for all instances. The active dataset will be up to 100 GB in size. A solutions architect has determined that peak operations will occur for 3 hours daily and will require a total of 225 MiBps of read throughput.

The solutions architect must design a Multi-AZ solution that makes a copy of the data available in another AWS Region for disaster recovery (DR). The DR copy has an RPO of less than 1 hour.

Which solution will meet these requirements?

A. Deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. Configure the file system for 75 MiBps of provisioned throughput. Implementreplication to a file system in the DR Region.

B. Deploy a new Amazon FSx for Lustre file system. Configure Bursting Throughput mode for the file system. Use AWS Backup to back up the file system to the DR Region.

C. Deploy a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput. Enable Multi-Attach for the EBSvolume. Use AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region.

D. Deploy an Amazon FSx for OpenZFS file system in both the production Region and the DR Region.Create an AWS DataSync scheduled task to replicate thedata from the production file system to the DR file system every 10 minutes.

Answer: A

Explanation:

The company should deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. The company should configure the file system for 75 MiBps of provisioned throughput. The company should implement replication to a file system in the DR Region. This solution will meet the requirements because Amazon EFS is a serverless, fully elastic file storage service that lets you share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files1. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances. A Multi-AZ file system replicates data across multiple Availability Zones (AZs) within a Region, providing high availability and durability2. By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput. Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system's size or burst credit balance3. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of less than 1 hour.

The other options are not correct because:

Deploying a new Amazon FSx for Lustre file system would not provide a single location for updates to application data for all instances. Amazon FSx for Lustre is a fully managed service that provides cost- effective, high-performance storage for compute workloads. However, it does not support concurrent write access from multiple instances. Using AWS Backup to back up the file system to the DR Region would not provide real-time replication of data. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. However, it does not support continuous data replication or cross-Region disaster recovery.

Deploying a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput would not provide a single location for updates to application data for all instances.

Amazon EBS is a service that provides persistent block storage volumes for use with Amazon EC2 instances.

However, it does not support concurrent access from multiple instances, unless Multi-Attach is enabled.

Enabling Multi-Attach for the EBS volume would not provide Multi-AZ resilience or cross-Region replication. Multi-Attach is a feature that enables you to attach an EBS volume to multiple EC2 instances within the same Availability Zone. Using AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region would not provide real-time replication of data. AWS Elastic Disaster Recovery (AWS DRS) is a service that enables you

to orchestrate and automate disaster recovery workflows across AWS Regions. However, it does not support continuous data replication or sub-hour RPOs.
Deploying an Amazon FSx for OpenZFS file system in both the production Region and the DR Region would not be as simple or cost-effective as using Amazon EFS. Amazon FSx for OpenZFS is a fully managed service that provides high-performance storage with strong data consistency and advanced data management features for Linux workloads. However, it requires more configuration and management than Amazon EFS, which is serverless and fully elastic. Creating an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes would not provide real-time replication of data.
AWS DataSync is a service that enables you to transfer data between on-premises storage and AWS services, or between AWS services. However, it does not support continuous data replication or sub-minute RPOs.
References:
https://aws.amazon.com/efs/
https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-azs
https://docs.aws.amazon.com/efs/latest/ug/performance.html#provisioned-throughput
https://docs.aws.amazon.com/efs/latest/ug/replication.html
https://aws.amazon.com/fsx/lustre/
https://aws.amazon.com/backup/
https://aws.amazon.com/ebs/
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html

QUESTION NO: 33
A company's web application uses an Amazon API Gateway API, AWS Lambda functions, and Amazon DynamoDB global tables to handle backend requests. The web application is deployed in two AWS Regions in an active-passive model. The company uses Amazon Route 53 for DNS. The web application requires a manual DNS update to fail over to the secondary Region. An analytics Lambda function runs in the same AWS account. The function has caused Lambda concurrency to reach 90% of the current quota on an average day. A recent surge in traffic for the analytics workload resulted in throttled Lambda requests and a poor user experience for the web application users. A solutions architect must increase the reliability of the web application. The solution must use an Amazon CloudWatch alarm to send an Amazon SNS notification when the Lambda concurrency reaches a specific utilization threshold. Which solution will meet these requirements with the LEAST operational overhead?

A. Set reserved concurrency on the web application Lambda functions. Implement Route 53 health checks and failover records to route traffic to the secondary Region. Configure the CloudWatch alarm to use the AWS Trusted Advisor ServiceLimitUsage metric and to send the SNS notification.

B. Set reserved concurrency on the web application Lambda functions. Implement Route 53 health checks and latency records to route traffic to the secondary Region. Configure the CloudWatch alarm to use the AWS Trusted Advisor ServiceLimitUsage metric and to send an SNS notification.

C. Set provisioned concurrency on the web application Lambda functions. Implement Route

53 health checks and failover records to route traffic to the secondary Region. Configure the CloudWatch alarm to use the Lambda ConcurrentExecutions metric and to send an SNS notification.

D. Set provisioned concurrency on the web application Lambda functions. Implement Route 53 health checks and geolocation records to route traffic to the secondary Region. Configure the CloudWatch alarm to use the Lambda ProvisionedConcurrencyInvocations metric and to send an SNS notification.

Answer: C

Explanation:

The use of provisioned concurrency ensures the web application's Lambda functions have pre-initialized execution environments, removing cold start latency and maintaining performance during high-traffic periods.

Route 53 health checks and failover records automate DNS failover to the secondary Region, improving application availability and reliability.

The CloudWatch alarm is configured to monitor the Lambda ConcurrentExecutions metric and send an SNS notification if concurrency usage nears the limit, enabling quick operational response.

This approach minimizes manual management, ensuring reliability and performance during peak traffic while meeting best practices for AWS Lambda and Route 53 failover.

QUESTION NO: 34

A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand.

Which solutions meet these requirements? (Choose two.)

A. Create an Amazon API Gateway REST API. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.

B. Create an Amazon API Gateway HTTP API. Configure this API with direct integrations to Dynamo DB by using API Gateway's AWS integration type.

C. Create an Amazon API Gateway HTTP API. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.

D. Create an accelerator in AWS Global Accelerator. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.

E. Create a Network Load Balancer. Configure listener rules to forward requests to the appropriate AWS Lambda functions

Answer: A C

Explanation:

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

html

QUESTION NO: 35

A company is running a containerized workload on AWS. The workload consists of several data-processing services that run on a group of Amazon EC2 instances.

The company uploads new data to an Amazon S3 bucket every night. A cron job on each

EC2 instance starts the data processing every night. The amount of uploaded data varies. The data-processing tasks can take hours to finish running. After the data is processed, the services remain idle until the next processing window occurs the next night. The company needs a solution to modernize the architecture and reduce the operational overhead.
Which solution will meet these requirements?

A. Migrate the workload to AWS Lambda functions that run the container images. Configure an Amazon EventBridge rule to filter S3 events and invoke the Lambda functions when data is uploaded to the S3 bucket.

B. Migrate the workload to run as tasks in an Amazon ECS cluster that runs on AWS Fargate. Create an AWS Step Functions state machine to invoke the Fargate tasks. Configure S3 Event Notifications to invoke the state machine tasks when data is uploaded to the S3 bucket.

C. Migrate the workload to run as tasks in an Amazon ECS cluster that runs on AWS Fargate. Create an AWS Step Functions state machine to invoke the Fargate tasks. Configure an Amazon EventBridge rule to invoke the state machine when data is uploaded to the S3 bucket.

D. Migrate the workload to AWS Lambda functions by packaging the container images as Lambda layers.Configure S3 Event Notifications to invoke the Lambda functions when data is uploaded to the S3 bucket.

Answer: C

Explanation:

The workload is containerized, runs for hours, and is event-driven by nightly data arrival in Amazon S3. The current architecture uses EC2 instances and cron jobs, which results in operational overhead (managing instances, patching, scaling, scheduling) and idle compute between processing windows.

A key constraint is that the processing tasks can take hours. AWS Lambda has maximum execution duration limits that make it unsuitable for multi-hour batch processing. Even though Lambda can run container images, it still must complete within Lambda's runtime limit. Packaging container images as Lambda layers is also not an appropriate pattern for long-running container workloads and adds complexity.

A modern, low-ops approach for long-running, containerized batch jobs is to run containers on AWS Fargate.

Fargate removes the need to manage EC2 instances and allows tasks to run for extended periods as needed, scaling based on demand. Because the workload is composed of several data-processing services that likely need orchestration (for example, fan-out, sequencing, retries, parallelism), AWS Step Functions is well suited to coordinate the workflow and invoke the appropriate ECS tasks.

For triggering based on new S3 data, Amazon EventBridge provides a managed, scalable event bus for AWS service events, including S3 object events, and can route events to targets such as Step Functions state machines. Using EventBridge reduces the need for direct point-to-point notification wiring and provides centralized event routing, filtering, and monitoring.

Option C combines all the right elements: it runs the containers as ECS tasks on Fargate to eliminate EC2 management and idle capacity, uses Step Functions to orchestrate tasks that can run for hours, and uses EventBridge to trigger the state machine when new data is

uploaded to S3. This replaces the per-instance cron scheduling with an event-driven serverless orchestration model and significantly reduces operational overhead.
Option B is close but is less appropriate as written because S3 Event Notifications are typically configured to send to Amazon SQS, Amazon SNS, or AWS Lambda. Triggering Step Functions directly is more naturally handled through EventBridge rules. EventBridge is also the recommended event routing layer for integrating service events into workflows.
Option A is not suitable because Lambda is not designed for multi-hour processing jobs due to runtime limits.
Option D is incorrect because Lambda layers are for sharing libraries and runtime dependencies, not for packaging multi-hour container workloads. It also still depends on Lambda runtime limits and does not match the operational model for long-running batch processing.
Therefore, option C is the best modernization approach with the least operational overhead.
References:AWS documentation on AWS Fargate for running container workloads without managing EC2 instances and supporting long-running tasks.AWS documentation on AWS Step Functions for orchestrating long-running workflows, retries, parallelism, and service integrations including Amazon ECS.AWS documentation on Amazon EventBridge for routing Amazon S3 object events to targets such as Step Functions state machines for event-driven architectures.