

## [2025-November-NewBraindump2go SAA-C03 VCE Questions Free][Q976-Q1010]

2025/November Latest Braindump2go SAA-C03 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go SAA-C03 Real Exam Questions!QUESTION 976A company uses Amazon S3 to host its static website. The

company wants to add a contact form to the webpage. The contact form will have dynamic server-side components for users to input their name, email address, phone number, and user message. The company expects fewer than 100 site visits each month. The contact form must notify the company by email when a customer fills out the form. Which solution will meet these requirements MOST cost-effectively?A. Host the dynamic contact form in Amazon Elastic Container Service (Amazon ECS). Set up Amazon Simple Email Service (Amazon SES) to connect to a third-party email provider.B. Create an Amazon API Gateway endpoint that returns the contact form from an AWS Lambda function. Configure another Lambda function on the API Gateway to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.C. Host the website by using AWS Amplify Hosting for static content and dynamic content. Use server-side scripting to build the contact form. Configure Amazon Simple Queue Service (Amazon SQS) to deliver the message to the company.D. Migrate the website from Amazon S3 to Amazon EC2 instances that run Windows Server. Use Internet Information Services (IIS) for Windows Server to host the webpage. Use client-side scripting to build the contact form. Integrate the form with Amazon WorkMail. Answer: BExplanation: Using API Gateway and Lambda enables serverless handling of form submissions with minimal cost and infrastructure. When coupled with Amazon SNS, it allows instant email notifications without running servers, making it ideal for low-traffic workloads.QUESTION 977A company creates dedicated AWS accounts in AWS Organizations for its business units. Recently, an important notification was sent to the root user email address of a business unit account instead of the assigned account owner. The company wants to ensure that all future notifications can be sent to different employees based on the notification categories of billing, operations, or security. Which solution will meet these requirements MOST securely?A. Configure each AWS account to use a single email address that the company manages. Ensure that all account owners can access the email account to receive notifications. Configure alternate contacts for each AWS account with corresponding distribution lists for the billing team, the security team, and the operations team for each business unit.B.

Configure each AWS account to use a different email distribution list for each business unit that the company manages. Configure each distribution list with administrator email addresses that can respond to alerts. Configure alternate contacts for each AWS account with corresponding distribution lists for the billing team, the security team, and the operations team for each business unit.C.

Configure each AWS account root user email address to be the individual company managed email address of one person from each business unit. Configure alternate contacts for each AWS account with corresponding distribution lists for the billing team, the security team, and the operations team for each business unit.D. Configure each AWS account root user to use email aliases that go to a centralized mailbox. Configure alternate contacts for each account by using a single business managed email distribution list each for the billing team, the security team, and the operations team. Answer: BQUESTION 978A company runs an ecommerce application on AWS. Amazon EC2 instances process purchases and store the purchase details in an Amazon Aurora PostgreSQL DB cluster. Customers are experiencing application timeouts during times of peak usage. A solutions architect needs to rearchitect the application so that the application can scale to meet peak usage demands. Which combination of actions will meet these requirements MOST cost-effectively? (Choose two.)A. Configure an Auto Scaling group of new EC2 instances to retry the purchases until the processing is complete. Update the applications to connect to the DB cluster by using Amazon RDS Proxy.B. Configure the application to use an Amazon ElastiCache cluster in front of the Aurora PostgreSQL DB cluster.C. Update the application to send the purchase requests to an Amazon Simple Queue Service (Amazon SQS) queue. Configure an Auto Scaling group of new EC2 instances that read from the SQS queue.D. Configure an AWS Lambda function to retry the ticket purchases until the processing is complete.E. Configure an Amazon API Gateway REST API with a usage plan. Answer: ACQUESTION 979A company that uses AWS Organizations runs 150 applications across 30 different AWS accounts. The company used AWS Cost and Usage Report to create a new report in the management account. The report is delivered to an Amazon S3 bucket that is replicated to a bucket in the data collection account. The company's senior leadership wants to view a custom dashboard that provides NAT gateway costs each day starting at the beginning of the current month. Which solution will meet these requirements?A. Share an Amazon QuickSight dashboard that includes the requested table visual. Configure QuickSight to use AWS DataSync to query the new report.B. Share an Amazon QuickSight dashboard that includes the requested table visual. Configure QuickSight to use Amazon Athena to query the new report.C. Share an Amazon CloudWatch dashboard that includes the requested table visual. Configure CloudWatch to use AWS DataSync to query the new report.D. Share an Amazon CloudWatch dashboard that includes the requested table visual. Configure CloudWatch to use Amazon Athena to query the new report. Answer: BQUESTION 980A company is hosting a

high-traffic static website on Amazon S3 with an Amazon CloudFront distribution that has a default TTL of 0 seconds. The company wants to implement caching to improve performance for the website. However, the company also wants to ensure that stale content is not served for more than a few minutes after a deployment. Which combination of caching methods should a solutions architect implement to meet these requirements? (Choose two.)

A. Set the CloudFront default TTL to 2 minutes.  
 B. Set a default TTL of 2 minutes on the S3 bucket.  
 C. Add a Cache-Control private directive to the objects in Amazon S3.  
 D. Create an AWS Lambda@Edge function to add an Expires header to HTTP responses. Configure the function to run on viewer response.  
 E. Add a Cache-Control max-age directive of 24 hours to the objects in Amazon S3. On deployment, create a CloudFront invalidation to clear any changed files from edge caches.

Answer: AC

QUESTION 981 A company runs its application by using Amazon EC2 instances and AWS Lambda functions. The EC2 instances run in private subnets of a VPC. The Lambda functions need direct network access to the EC2 instances for the application to work. The application will run for 1 year. The number of Lambda functions that the application uses will increase during the 1-year period. The company must minimize costs on all application resources. Which solution will meet these requirements?

A. Purchase an EC2 Instance Savings Plan. Connect the Lambda functions to the private subnets that contain the EC2 instances.  
 B. Purchase an EC2 Instance Savings Plan. Connect the Lambda functions to new public subnets in the same VPC where the EC2 instances run.  
 C. Purchase a Compute Savings Plan. Connect the Lambda functions to the private subnets that contain the EC2 instances.  
 D. Purchase a Compute Savings Plan. Keep the Lambda functions in the Lambda service VPC.

Answer: C

Explanation: Compute Savings Plan: This plan offers significant discounts on Lambda functions compared to on-demand pricing. Since the application will run for a year, a sustained use discount like Compute Savings Plan is ideal.

Private Subnets: Lambda functions in private subnets can directly access EC2 instances within the VPC without needing internet access, reducing security risks and potential egress costs.

QUESTION 982 A company has deployed a multi-account strategy on AWS by using AWS Control Tower. The company has provided individual AWS accounts to each of its developers. The company wants to implement controls to limit AWS resource costs that the developers incur. Which solution will meet these requirements with the LEAST operational overhead?

A. Instruct each developer to tag all their resources with a tag that has a key of CostCenter and a value of the developer's name. Use the required-tags AWS Config managed rule to check for the tag. Create an AWS Lambda function to terminate resources that do not have the tag. Configure AWS Cost Explorer to send a daily report to each developer to monitor their spending.  
 B. Use AWS Budgets to establish budgets for each developer account. Set up budget alerts for actual and forecast values to notify developers when they exceed or expect to exceed their assigned budget. Use AWS Budgets actions to apply a DenyAll policy to the developer's IAM role to prevent additional resources from being launched when the assigned budget is reached.  
 C. Use AWS Cost Explorer to monitor and report on costs for each developer account. Configure Cost Explorer to send a daily report to each developer to monitor their spending. Use AWS Cost Anomaly Detection to detect anomalous spending and provide alerts.  
 D. Use AWS Service Catalog to allow developers to launch resources within a limited cost range. Create AWS Lambda functions in each AWS account to stop running resources at the end of each work day. Configure the Lambda functions to resume the resources at the start of each work day.

Answer: B

Explanation: Taking into consideration that AWS Budgets is allowing to will inform you that you exceeded budget and execute actions like for example IAM actions to prevent running new resources in cloud, I think this option is a good and resonable move. In case of need budget can be always increased and "chains" disabled.

<https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-controls.html>

QUESTION 983 A solutions architect is designing a three-tier web application. The architecture consists of an internet-facing Application Load Balancer (ALB) and a web tier that is hosted on Amazon EC2 instances in private subnets. The application tier with the business logic runs on EC2 instances in private subnets. The database tier consists of Microsoft SQL Server that runs on EC2 instances in private subnets. Security is a high priority for the company. Which combination of security group configurations should the solutions architect use? (Choose three.)

A. Configure the security group for the web tier to allow inbound HTTPS traffic from the security group for the ALB.  
 B. Configure the security group for the web tier to allow outbound HTTPS traffic to 0.0.0.0/0.  
 C. Configure the security group for the database tier to allow inbound Microsoft SQL Server traffic from the security group for the application tier.  
 D. Configure the security group for the database tier to allow outbound HTTPS traffic and Microsoft SQL Server traffic to the security group for the web tier.  
 E. Configure the security group for the application tier to allow inbound HTTPS traffic from the security group for the web tier.  
 F. Configure the security group for the application tier to allow outbound HTTPS traffic and Microsoft SQL Server traffic to the security group for the web tier.

Answer: ACE

Explanation: Security Group is protecting instances, it's statefull. by defoult is allowing for outgoing traffic but not incomming.hence we need to allow for inboud traffic. path looks like below

ALB >> HTTPS >> WEB tier >> HTTPS >> Application >> SQL traffic >> SQL DB

hence we need allow for incoming https traffic on web tier then incoming http on app tier and on the end for incoming sql traffic on DB tier

QUESTION 984 A company has released a new version of its production application. The company's workload uses Amazon EC2, AWS Lambda, AWS Fargate, and Amazon SageMaker. The

company wants to cost optimize the workload now that usage is at a steady state. The company wants to cover the most services with the fewest savings plans. Which combination of savings plans will meet these requirements? (Choose two.)  
A. Purchase an EC2 Instance Savings Plan for Amazon EC2 and SageMaker.  
B. Purchase a Compute Savings Plan for Amazon EC2, Lambda, and SageMaker.  
C. Purchase a SageMaker Savings Plan.  
D. Purchase a Compute Savings Plan for Lambda, Fargate, and Amazon EC2.  
E. Purchase an EC2 Instance Savings Plan for Amazon EC2 and Fargate.

Answer: CDExplanation: <https://aws.amazon.com/savingsplans/ml-pricing/>

<https://aws.amazon.com/savingsplans/compute-pricing/> QUESTION 985A  
A company uses a Microsoft SQL Server database. The company's applications are connected to the database. The company wants to migrate to an Amazon Aurora PostgreSQL database with minimal changes to the application code. Which combination of steps will meet these requirements? (Choose two.)  
A. Use the AWS Schema Conversion Tool (AWS SCT) to rewrite the SQL queries in the applications.  
B. Enable Babelfish on Aurora PostgreSQL to run the SQL queries from the applications.  
C. Migrate the database schema and data by using the AWS Schema Conversion Tool (AWS SCT) and AWS Database Migration Service (AWS DMS).  
D. Use Amazon RDS Proxy to connect the applications to Aurora PostgreSQL.  
E. Use AWS Database Migration Service (AWS DMS) to rewrite the SQL queries in the applications.

Answer: BCExplanation: DMS will allow for DATABASE migration and use AWS Schema Conversion Tool (AWS SCT) to create some or all of the target tables, indexes, views, triggers, and so on.

<https://docs.aws.amazon.com/dms/latest/userguide/Welcome.html> To minimize amount of code which need to me changes we need to use babelfish  
<https://aws.amazon.com/rds/aurora/babelfish/> QUESTION 986A company plans to rehost an application to Amazon EC2 instances that use Amazon Elastic Block Store (Amazon EBS) as the attached storage. A solutions architect must design a solution to ensure that all newly created Amazon EBS volumes are encrypted by default. The solution must also prevent the creation of unencrypted EBS volumes. Which solution will meet these requirements?  
A. Configure the EC2 account attributes to always encrypt new EBS volumes.  
B. Use AWS Config. Configure the encrypted-volumes identifier. Apply the default AWS Key Management Service (AWS KMS) key.  
C. Configure AWS Systems Manager to create encrypted copies of the EBS volumes. Reconfigure the EC2 instances to use the encrypted volumes.  
D. Create a customer managed key in AWS Key Management Service (AWS KMS). Configure AWS Migration Hub to use the key when the company migrates workloads.

Answer: AExplanation: <https://repost.aws/knowledge-center/ebs-automatic-encryption> QUESTION 987An ecommerce company wants to collect user clickstream data from the company's website for real-time analysis. The website experiences fluctuating traffic patterns throughout the day. The company needs a scalable solution that can adapt to varying levels of traffic. Which solution will meet these requirements?  
A. Use a data stream in Amazon Kinesis Data Streams in on-demand mode to capture the clickstream data. Use AWS Lambda to process the data in real time.  
B. Use Amazon Kinesis Data Firehose to capture the clickstream data. Use AWS Glue to process the data in real time.  
C. Use Amazon Kinesis Video Streams to capture the clickstream data. Use AWS Glue to process the data in real time.  
D. Use Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) to capture the clickstream data. Use AWS Lambda to process the data in real time.

Answer: AExplanation: Apache Flink (previously known as Amazon Kinesis Data Analytics) seems to not allowing sent data directly to Lambda...Glue is allowing to integrate data from couple of sources in to one.  
<https://docs.aws.amazon.com/lambda/latest/dg/with-kinesis.html>

<https://aws.amazon.com/kinesis/data-streams/features/?nc=sn&loc=2> QUESTION 988A global company runs its workloads on AWS. The company's application uses Amazon S3 buckets across AWS Regions for sensitive data storage and analysis. The company stores millions of objects in multiple S3 buckets daily. The company wants to identify all S3 buckets that are not versioning-enabled. Which solution will meet these requirements?  
A. Set up an AWS CloudTrail event that has a rule to identify all S3 buckets that are not versioning-enabled across Regions.  
B. Use Amazon S3 Storage Lens to identify all S3 buckets that are not versioning-enabled across Regions.  
C. Enable IAM Access Analyzer for S3 to identify all S3 buckets that are not versioning-enabled across Regions.  
D. Create an S3 Multi-Region Access Point to identify all S3 buckets that are not versioning-enabled across Regions.

Answer: BExplanation: S3 Storage Lens "can also identify buckets that aren't following data-protection best practices, such as using S3 Replication or S3 Versioning. "

[https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage\\_lens\\_basics\\_metrics\\_recommendations.html](https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens_basics_metrics_recommendations.html) QUESTION 989A company needs to optimize its Amazon S3 storage costs for an application that generates many files that cannot be recreated. Each file is approximately 5 MB and is stored in Amazon S3 Standard storage. The company must store the files for 4 years before the files can be deleted. The files must be immediately accessible. The files are frequently accessed in the first 30 days of object creation, but they are rarely accessed after the first 30 days. Which solution will meet these requirements MOST cost-effectively?  
A. Create an S3 Lifecycle policy to move the files to S3 Glacier Instant Retrieval 30 days after object creation. Delete the files 4 years after object creation.  
B. Create an S3 Lifecycle policy to move the files to S3 One Zone-Infrequent Access (S3 One Zone-IA) 30 days after object creation. Delete the files 4 years after object creation.  
C. Create an S3 Lifecycle policy to move the files to S3

Standard-Infrequent Access (S3 Standard-IA) 30 days after object creation. Delete the files 4 years after object creation.D. Create an S3 Lifecycle policy to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) 30 days after object creation. Move the files to S3 Glacier Flexible Retrieval 4 years after object creation.Answer: CExplanation:Requirements:- frequently accessed for 30 days- lower costFeatures for S3 Standard-IA:- Infrequently accessed objects- Milliseconds to accessQUESTION 990A company runs its critical storage application in the AWS Cloud. The application uses Amazon S3 in two AWS Regions. The company wants the application to send remote user data to the nearest S3 bucket with no public network congestion. The company also wants the application to fail over with the least amount of management of Amazon S3.Which solution will meet these requirements?A. Implement an active-active design between the two Regions. Configure the application to use the regional S3 endpoints closest to the user.B. Use an active-passive configuration with S3 Multi-Region Access Points. Create a global endpoint for each of the Regions.C. Send user data to the regional S3 endpoints closest to the user. Configure an S3 cross-account replication rule to keep the S3 buckets synchronized.D. Set up Amazon S3 to use Multi-Region Access Points in an active-active configuration with a single global endpoint. Configure S3 Cross-Region Replication.Answer: DExplanation:Using a Multi-region Accesspoint in an Active-Active setup will send data to the closest Region, without accessing the internet: "send remote user data to the nearest S3 bucket with no public network congestion"<https://docs.aws.amazon.com/AmazonS3/latest/userguide/MultiRegionAccessPoints.html>QUESTION 991A company is migrating a data center from its on-premises location to AWS. The company has several legacy applications that are hosted on individual virtual servers. Changes to the application designs cannot be made. Each individual virtual server currently runs as its own EC2 instance. A solutions architect needs to ensure that the applications are reliable and fault tolerant after migration to AWS. The applications will run on Amazon EC2 instances.Which solution will meet these requirements?A. Create an Auto Scaling group that has a minimum of one and a maximum of one. Create an Amazon Machine Image (AMI) of each application instance. Use the AMI to create EC2 instances in the Auto Scaling group. Configure an Application Load Balancer in front of the Auto Scaling group.B. Use AWS Backup to create an hourly backup of the EC2 instance that hosts each application. Store the backup in Amazon S3 in a separate Availability Zone. Configure a disaster recovery process to restore the EC2 instance for each application from its most recent backup.C. Create an Amazon Machine Image (AMI) of each application instance. Launch two new EC2 instances from the AMI. Place each EC2 instance in a separate Availability Zone. Configure a Network Load Balancer that has the EC2 instances as targets.D. Use AWS Mitigation Hub Refactor Spaces to migrate each application off the EC2 instance. Break down functionality from each application into individual components. Host each application on Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch type.Answer: CQUESTION 992A company wants to isolate its workloads by creating an AWS account for each workload. The company needs a solution that centrally manages networking components for the workloads. The solution also must create accounts with automatic security controls (guardrails). Which solution will meet these requirements with the LEAST operational overhead?A. Use AWS Control Tower to deploy accounts. Create a networking account that has a VPC with private subnets and public subnets. Use AWS Resource Access Manager (AWS RAM) to share the subnets with the workload accounts.B. Use AWS Organizations to deploy accounts. Create a networking account that has a VPC with private subnets and public subnets. Use AWS Resource Access Manager (AWS RAM) to share the subnets with the workload accounts.C. Use AWS Control Tower to deploy accounts. Deploy a VPC in each workload account. Configure each VPC to route through an inspection VPC by using a transit gateway attachment.D. Use AWS Organizations to deploy accounts. Deploy a VPC in each workload account. Configure each VPC to route through an inspection VPC by using a transit gateway attachment.Answer: AExplanation:AWS Control Tower provides a pre-packaged set of guardrails (policies) and blueprints (best-practice configurations) to ensure that the environment complies with security and compliance standards. It's designed to simplify the process of creating and managing a multi-account AWS environment while maintaining security and compliance.<https://docs.aws.amazon.com/controlltower/latest/userguide/what-is-control-tower.html>QUESTION 993A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The website serves static content. Website traffic is increasing. The company wants to minimize the website hosting costs.Which solution will meet these requirements?A. Move the website to an Amazon S3 bucket. Configure an Amazon CloudFront distribution for the S3 bucket.B. Move the website to an Amazon S3 bucket. Configure an Amazon ElastiCache cluster for the S3 bucket.C. Move the website to AWS Amplify. Configure an ALB to resolve to the Amplify website.D. Move the website to AWS Amplify. Configure EC2 instances to cache the website.Answer: AExplanation:Amazon CloudFront:Uses the durable storage of Amazon Simple Storage Service (Amazon S3) - This solution creates an Amazon S3 bucket to host your static website's content. To update your website, just upload your new files to the S3 bucket.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/getting-started-secure-static-website-cloudformation-template.html>QUESTION 994A company is implementing a shared storage solution for a media application that the company hosts on

AWS. The company needs the ability to use SMB clients to access stored data. Which solution will meet these requirements with the LEAST administrative overhead? A. Create an AWS Storage Gateway Volume Gateway. Create a file share that uses the required client protocol. Connect the application server to the file share. B. Create an AWS Storage Gateway Tape Gateway. Configure tapes to use Amazon S3. Connect the application server to the Tape Gateway. C. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance. Connect the application server to the file share. D. Create an Amazon FSx for Windows File Server file system. Connect the application server to the file system.

Answer: DExplanation: <https://aws.amazon.com/fsx/windows/> QUESTION 995 A company is designing its production application's disaster recovery (DR) strategy. The application is backed by a MySQL database on an Amazon Aurora cluster in the us-east-1 Region. The company has chosen the us-west-1 Region as its DR Region. The company's target recovery point objective (RPO) is 5 minutes and the target recovery time objective (RTO) is 20 minutes. The company wants to minimize configuration changes. Which solution will meet these requirements with the MOST operational efficiency? A. Create an Aurora read replica in us-west-1 similar in size to the production application's Aurora MySQL cluster writer instance. B. Convert the Aurora cluster to an Aurora global database.

Configure managed failover. C. Create a new Aurora cluster in us-west-1 that has Cross-Region Replication. D. Create a new Aurora cluster in us-west-1. Use AWS Database Migration Service (AWS DMS) to sync both clusters.

Answer: BExplanation: Cross-Region disaster recovery If your primary Region suffers a performance degradation or outage, you can promote one of the secondary Regions to take read/write responsibilities. An Aurora cluster can recover in less than 1 minute, even in the event of a complete Regional outage. This provides your application with an effective recovery point objective (RPO) of 1 second and a recovery time objective (RTO) of less than 1 minute, providing a strong foundation for a global business continuity plan.

<https://aws.amazon.com/rds/aurora/global-database/> QUESTION 996 A company runs a critical data analysis job each week before the first day of the work week. The job requires at least 1 hour to complete the analysis. The job is stateful and cannot tolerate interruptions. The company needs a solution to run the job on AWS. Which solution will meet these requirements? A. Create a container for the job. Schedule the job to run as an AWS Fargate task on an Amazon Elastic Container Service (Amazon ECS) cluster by using Amazon EventBridge Scheduler. B. Configure the job to run in an AWS Lambda function. Create a scheduled rule in Amazon EventBridge to invoke the Lambda function. C. Configure an Auto Scaling group of Amazon EC2 Spot Instances that run Amazon Linux. Configure a crontab entry on the instances to run the analysis. D. Configure an AWS DataSync task to run the job. Configure a cron expression to run the task on a schedule.

Answer: AQUESTION 997 A company runs workloads in the AWS Cloud. The company wants to centrally collect security data to assess security across the entire company and to improve workload protection. Which solution will meet these requirements with the LEAST development effort? A. Configure a data lake in AWS Lake Formation. Use AWS Glue crawlers to ingest the security data into the data lake. B. Configure an AWS Lambda function to collect the security data in .csv format. Upload the data to an Amazon S3 bucket. C. Configure a data lake in Amazon Security Lake to collect the security data. Upload the data to an Amazon S3 bucket. D. Configure an AWS Database Migration Service (AWS DMS) replication instance to load the security data into an Amazon RDS cluster.

Answer: CExplanation: Amazon Security Lake automatically centralizes security data from AWS environments, you can get a more complete understanding of your security data across your entire organization. You can also improve the protection. QUESTION 998 A company is migrating five on-premises applications to VPCs in the AWS Cloud. Each application is currently deployed in isolated virtual networks on premises and should be deployed similarly in the AWS Cloud. The applications need to reach a shared services VPC. All the applications must be able to communicate with each other. If the migration is successful, the company will repeat the migration process for more than 100 applications. Which solution will meet these requirements with the LEAST administrative overhead? A. Deploy software VPN tunnels between the application VPCs and the shared services VPC. Add routes between the application VPCs in their subnets to the shared services VPC. B. Deploy VPC peering connections between the application VPCs and the shared services VPC. Add routes between the application VPCs in their subnets to the shared services VPC through the peering connection. C. Deploy an AWS Direct Connect connection between the application VPCs and the shared services VPC. Add routes from the application VPCs in their subnets to the shared services VPC and the applications VPCs. Add routes from the shared services VPC subnets to the applications VPCs. D. Deploy a transit gateway with associations between the transit gateway and the application VPCs and the shared services VPC. Add routes between the application VPCs in their subnets and the application VPCs to the shared services VPC through the transit gateway.

Answer: DExplanation: It will allow for inter-VPC communication for all 5 applications/VPC, reach shared resource/VPC and in the future it will be easy to allow for inter-communication between even 100 VPCs (applications).

<https://aws.amazon.com/transit-gateway/> QUESTION 999 A company wants to use Amazon Elastic Container Service (Amazon ECS) to run its on-premises application in a hybrid environment. The application currently runs on containers on premises. The company needs a single container solution that can scale in an on-premises, hybrid, or cloud environment. The company must run

new application containers in the AWS Cloud and must use a load balancer for HTTP traffic. Which combination of actions will meet these requirements? (Choose two.)  
A. Set up an ECS cluster that uses the AWS Fargate launch type for the cloud application containers. Use an Amazon ECS Anywhere external launch type for the on-premises application containers.  
B. Set up an Application Load Balancer for cloud ECS services.  
C. Set up a Network Load Balancer for cloud ECS services.  
D. Set up an ECS cluster that uses the AWS Fargate launch type. Use Fargate for the cloud application containers and the on-premises application containers.  
E. Set up an ECS cluster that uses the Amazon EC2 launch type for the cloud application containers. Use Amazon ECS Anywhere with an AWS Fargate launch type for the on-premises application containers.

Answer: AB

Explanation: We need to load-balance HTTP traffic hence Application Load Balancer is needed. Because Customer want to use container solution we need to use ECS with Fargate which will launch cloud applications. To run on-premises applications in containers we need to use ECS Anywhere. [https://docs.aws.amazon.com/AmazonECS/latest/developerguide/AWS\\_Fargate.html](https://docs.aws.amazon.com/AmazonECS/latest/developerguide/AWS_Fargate.html)

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html>

QUESTION 1000 A company is migrating its workloads to AWS. The company has sensitive and critical data in on-premises relational databases that run on SQL Server instances. The company wants to use the AWS Cloud to increase security and reduce operational overhead for the databases. Which solution will meet these requirements?  
A. Migrate the databases to Amazon EC2 instances. Use an AWS Key Management Service (AWS KMS) AWS managed key for encryption.  
B. Migrate the databases to a Multi-AZ Amazon RDS for SQL Server DB instance. Use an AWS Key Management Service (AWS KMS) AWS managed key for encryption.  
C. Migrate the data to an Amazon S3 bucket. Use Amazon Macie to ensure data security.  
D. Migrate the databases to an Amazon DynamoDB table. Use Amazon CloudWatch Logs to ensure data security.

Answer: B

QUESTION 1001 A company wants to migrate an application to AWS. The company wants to increase the application's current availability. The company wants to use AWS WAF in the application's architecture. Which solution will meet these requirements?  
A. Create an Auto Scaling group that contains multiple Amazon EC2 instances that host the application across two Availability Zones. Configure an Application Load Balancer (ALB) and set the Auto Scaling group as the target. Connect a WAF to the ALB.  
B. Create a cluster placement group that contains multiple Amazon EC2 instances that hosts the application. Configure an Application Load Balancer and set the EC2 instances as the targets. Connect a WAF to the placement group.  
C. Create two Amazon EC2 instances that host the application across two Availability Zones. Configure the EC2 instances as the targets of an Application Load Balancer (ALB). Connect a WAF to the ALB.  
D. Create an Auto Scaling group that contains multiple Amazon EC2 instances that host the application across two Availability Zones. Configure an Application Load Balancer (ALB) and set the Auto Scaling group as the target. Connect a WAF to the Auto Scaling group.

Answer: A

QUESTION 1002 A company manages a data lake in an Amazon S3 bucket that numerous applications access. The S3 bucket contains a unique prefix for each application. The company wants to restrict each application to its specific prefix and to have granular control of the objects under each prefix. Which solution will meet these requirements with the LEAST operational overhead?  
A. Create dedicated S3 access points and access point policies for each application.  
B. Create an S3 Batch Operations job to set the ACL permissions for each object in the S3 bucket.  
C. Replicate the objects in the S3 bucket to new S3 buckets for each application. Create replication rules by prefix.  
D. Replicate the objects in the S3 bucket to new S3 buckets for each application. Create dedicated S3 access points for each application.

Answer: A

Explanation: When using SQS we will be sure that all images will be processed and hence to process we need 2 min and 512 MB of memory (Lambda is allowing upto 15 min and upto 10K MB) Lambda should be perfect scalable solution which will allow for almost in real time image processing.

QUESTION 1004 A company wants to improve the availability and performance of its hybrid application. The application consists of a stateful TCP-based workload hosted on Amazon EC2 instances

in different AWS Regions and a stateless UDP-based workload hosted on premises. Which combination of actions should a solutions architect take to improve availability and performance? (Choose two.)

A. Create an accelerator using AWS Global Accelerator.  
Add the load balancers as endpoints.

B. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the load balancers.

C. Configure two Application Load Balancers in each Region. The first will route to the EC2 endpoints, and the second will route to the on-premises endpoints.

D. Configure a Network Load Balancer in each Region to address the EC2 endpoints. Configure a Network Load Balancer in each Region that routes to the on-premises endpoints.

E. Configure a Network Load Balancer in each Region to address the EC2 endpoints. Configure an Application Load Balancer in each Region that routes to the on-premises endpoints.

Answer: ADQUESTION 1005A company runs a self-managed Microsoft SQL Server on Amazon EC2 instances and Amazon Elastic Block Store (Amazon EBS). Daily snapshots are taken of the EBS volumes. Recently, all the company's EBS snapshots were accidentally deleted while running a snapshot cleaning script that deletes all expired EBS snapshots. A solutions architect needs to update the architecture to prevent data loss without retaining EBS snapshots indefinitely. Which solution will meet these requirements with the LEAST development effort?

A. Change the IAM policy of the user to deny EBS snapshot deletion.

B. Copy the EBS snapshots to another AWS Region after completing the snapshots daily.

C. Create a 7-day EBS snapshot retention rule in Recycle Bin and apply the rule for all snapshots.

D. Copy EBS snapshots to Amazon S3 Standard-Infrequent Access (S3 Standard-IA).

Answer: CExplanation:

<https://aws.amazon.com/blogs/aws/new-recycle-bin-for-ebs-snapshots/>

QUESTION 1006A company wants to use an AWS CloudFormation stack for its application in a test environment. The company stores the CloudFormation template in an Amazon S3 bucket that blocks public access. The company wants to grant CloudFormation access to the template in the S3 bucket based on specific user requests to create the test environment. The solution must follow security best practices. Which solution will meet these requirements?

A. Create a gateway VPC endpoint for Amazon S3. Configure the CloudFormation stack to use the S3 object URL.

B. Create an Amazon API Gateway REST API that has the S3 bucket as the target. Configure the CloudFormation stack to use the API Gateway URL.

C. Create a presigned URL for the template object. Configure the CloudFormation stack to use the presigned URL.

D. Allow public access to the template object in the S3 bucket. Block the public access after the test environment is created.

Answer: CExplanation:<https://docs.aws.amazon.com/AmazonS3/latest/userguide/ShareObjectPreSignedURL.html>

QUESTION 1007A company has applications that run in an organization in AWS Organizations. The company outsources operational support of the applications. The company needs to provide access for the external support engineers without compromising security. The external support engineers need access to the AWS Management Console. The external support engineers also need operating system access to the company's fleet of Amazon EC2 instances that run Amazon Linux in private subnets. Which solution will meet these requirements MOST securely?

A. Confirm that AWS Systems Manager Agent (SSM Agent) is installed on all instances. Assign an instance profile with the necessary policy to connect to Systems Manager. Use AWS IAM Identity Center to provide the external support engineers console access. Use Systems Manager Session Manager to assign the required permissions.

B. Confirm that AWS Systems Manager Agent (SSM Agent) is installed on all instances. Assign an instance profile with the necessary policy to connect to Systems Manager. Use Systems Manager Session Manager to provide local IAM user credentials in each AWS account to the external support engineers for console access.

C. Confirm that all instances have a security group that allows SSH access only from the external support engineers' source IP address ranges. Provide local IAM user credentials in each AWS account to the external support engineers for console access. Provide each external support engineer an SSH key pair to log in to the application instances.

D. Create a bastion host in a public subnet. Set up the bastion host security group to allow access from only the external engineers' IP address ranges. Ensure that all instances have a security group that allows SSH access from the bastion host. Provide each external support engineer an SSH key pair to log in to the application instances. Provide local account IAM user credentials to the engineers for console access.

Answer: AExplanation: Systems Manager Session Manager allows secure, auditable, and controlled access to your EC2 instances without needing to open SSH ports or manage SSH keys, reducing the attack surface. Local IAM user credentials are less secure and harder to manage at scale compared to using IAM Identity Center.

QUESTION 1008A company uses Amazon RDS for PostgreSQL to run its applications in the us-east-1 Region. The company also uses machine learning (ML) models to forecast annual revenue based on near real-time reports. The reports are generated by using the same RDS for PostgreSQL database. The database performance slows during business hours. The company needs to improve database performance. Which solution will meet these requirements MOST cost-effectively?

A. Create a cross-Region read replica. Configure the reports to be generated from the read replica.

B. Activate Multi-AZ DB instance deployment for RDS for PostgreSQL. Configure the reports to be generated from the standby database.

C. Use AWS Data Migration Service (AWS DMS) to logically replicate data to a new database. Configure the reports to be generated from the new database.

D. Create a read replica in us-east-1. Configure the reports to be generated from the read replica.

Answer: DExplanation:

Read replicas are typically less expensive than setting up a cross-Region replica or activating Multi-AZ deployments. You only pay for the additional read replica, without the overhead costs associated with cross-Region data transfer or maintaining a synchronous standby in Multi-AZ setups.

QUESTION 1009

A company hosts its multi-tier, public web application in the AWS Cloud. The web application runs on Amazon EC2 instances, and its database runs on Amazon RDS. The company is anticipating a large increase in sales during an upcoming holiday weekend. A solutions architect needs to build a solution to analyze the performance of the web application with a granularity of no more than 2 minutes. What should the solutions architect do to meet this requirement?

A. Send Amazon CloudWatch logs to Amazon Redshift. Use Amazon QuickSight to perform further analysis.

B. Enable detailed monitoring on all EC2 instances. Use Amazon CloudWatch metrics to perform further analysis.

C. Create an AWS Lambda function to fetch EC2 logs from Amazon CloudWatch Logs. Use Amazon CloudWatch metrics to perform further analysis.

D. Send EC2 logs to Amazon S3. Use Amazon Redshift to fetch logs from the S3 bucket to process raw data for further analysis with Amazon QuickSight.

Answer: B

Explanation: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-cloudwatch-new.html>

QUESTION 1010

A company runs an application that stores and shares photos. Users upload the photos to an Amazon S3 bucket. Every day, users upload approximately 150 photos. The company wants to design a solution that creates a thumbnail of each new photo and stores the thumbnail in a second S3 bucket. Which solution will meet these requirements MOST cost-effectively?

A. Configure an Amazon EventBridge scheduled rule to invoke a script every minute on a long-running Amazon EMR cluster. Configure the script to generate thumbnails for the photos that do not have thumbnails. Configure the script to upload the thumbnails to the second S3 bucket.

B. Configure an Amazon EventBridge scheduled rule to invoke a script every minute on a memory-optimized Amazon EC2 instance that is always on. Configure the script to generate thumbnails for the photos that do not have thumbnails. Configure the script to upload the thumbnails to the second S3 bucket.

C. Configure an S3 event notification to invoke an AWS Lambda function each time a user uploads a new photo to the application. Configure the Lambda function to generate a thumbnail and to upload the thumbnail to the second S3 bucket.

D. Configure S3 Storage Lens to invoke an AWS Lambda function each time a user uploads a new photo to the application. Configure the Lambda function to generate a thumbnail and to upload the thumbnail to a second S3 bucket.

Answer: C

QUESTION 1011

A company has stored millions of objects across multiple prefixes in an Amazon S3 bucket by using the Amazon S3 Glacier Deep Archive storage class. The company needs to delete all data older than 3 years except for a subset of data that must be retained. The company has identified the data that must be retained and wants to implement a serverless solution. Which solution will meet these requirements?

A. Use S3 Inventory to list all objects. Use the AWS CLI to create a script that runs on an Amazon EC2 instance that deletes objects from the inventory list.

B. Use AWS Batch to delete objects older than 3 years except for the data that must be retained.

C. Provision an AWS Glue crawler to query objects older than 3 years. Save the manifest file of old objects. Create a script to delete objects in the manifest.

D. Enable S3 Inventory. Create an AWS Lambda function to filter and delete objects. Invoke the Lambda function with S3 Batch Operations to delete objects by using the inventory reports.

Answer: D

QUESTION 1012

A company is building an application on AWS. The application uses multiple AWS Lambda functions to retrieve sensitive data from a single Amazon S3 bucket for processing. The company must ensure that only authorized Lambda functions can access the data. The solution must comply with the principle of least privilege. Which solution will meet these requirements?

A. Grant full S3 bucket access to all Lambda functions through a shared IAM role.

B. Configure the Lambda functions to run within a VPC. Configure a bucket policy to grant access based on the Lambda functions' VPC endpoint IP addresses.

C. Create individual IAM roles for each Lambda function. Grant the IAM roles access to the S3 bucket. Assign each IAM role as the Lambda execution role for its corresponding Lambda function.

D. Configure a bucket policy granting access to the Lambda functions based on their function ARNs.

Answer: C

QUESTION 1013

A company has developed a non-production application that is composed of multiple microservices for each of the company's business units. A single development team maintains all the microservices. The current architecture uses a static web frontend and a Java-based backend that contains the application logic. The architecture also uses a MySQL database that the company hosts on an Amazon EC2 instance. The company needs to ensure that the application is secure and available globally. Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon CloudFront and AWS Amplify to host the static web frontend. Refactor the microservices to use AWS Lambda functions that the microservices access by using Amazon API Gateway. Migrate the MySQL database to an Amazon EC2 Reserved Instance.

B. Use Amazon CloudFront and Amazon S3 to host the static web frontend. Refactor the microservices to use AWS Lambda functions that the microservices access by using Amazon API Gateway. Migrate the MySQL database to Amazon RDS for MySQL.

C. Use Amazon CloudFront and Amazon S3 to host the static web frontend. Refactor the microservices to use AWS Lambda functions that are in a target group behind a Network Load Balancer. Migrate the MySQL database to Amazon RDS for MySQL.

D. Use Amazon S3 to host the static web frontend. Refactor the microservices to use AWS Lambda functions that are in a target group behind an Application Load Balancer. Migrate the

MySQL database to an Amazon EC2 Reserved Instance. Answer: B  
QUESTION 1014A video game company is deploying a new gaming application to its global users. The company requires a solution that will provide near real-time reviews and rankings of the players. A solutions architect must design a solution to provide fast access to the data. The solution must also ensure the data persists on disks in the event that the company restarts the application. Which solution will meet these requirements with the LEAST operational overhead?  
A. Configure an Amazon CloudFront distribution with an Amazon S3 bucket as the origin. Store the player data in the S3 bucket.  
B. Create Amazon EC2 instances in multiple AWS Regions. Store the player data on the EC2 instances. Configure Amazon Route 53 with geolocation records to direct users to the closest EC2 instance.  
C. Deploy an Amazon ElastiCache for Redis cluster. Store the player data in the ElastiCache cluster.  
D. Deploy an Amazon ElastiCache for Memcached cluster. Store the player data in the ElastiCache cluster.

Answer: C  
Explanation: Amazon ElastiCache for Redis provides in-memory caching which ensures low latency and high throughput, perfect for near real-time access to player reviews and rankings. Redis supports data persistence by snapshotting data to disk (RDB snapshots) and appending changes to a log (AOF), ensuring that the data is not lost even if the application restarts.

QUESTION 1015A company is designing an application on AWS that processes sensitive data. The application stores and processes financial data for multiple customers. To meet compliance requirements, the data for each customer must be encrypted separately at rest by using a secure, centralized key management solution. The company wants to use AWS Key Management Service (AWS KMS) to implement encryption. Which solution will meet these requirements with the LEAST operational overhead?  
A. Generate a unique encryption key for each customer. Store the keys in an Amazon S3 bucket. Enable server-side encryption.  
B. Deploy a hardware security appliance in the AWS environment that securely stores customer-provided encryption keys. Integrate the security appliance with AWS KMS to encrypt the sensitive data in the application.  
C. Create a single AWS KMS key to encrypt all sensitive data across the application.  
D. Create separate AWS KMS keys for each customer's data that have granular access control and logging enabled.

Answer: D  
Explanation: While enabling server-side encryption in S3 can manage encryption, it does not offer the same level of control and auditing as AWS KMS. Managing individual keys manually in S3 would also increase operational overhead.

QUESTION 1016A company needs to design a resilient web application to process customer orders. The web application must automatically handle increases in web traffic and application usage without affecting the customer experience or losing customer orders. Which solution will meet these requirements?  
A. Use a NAT gateway to manage web traffic. Use Amazon EC2 Auto Scaling groups to receive, process, and store processed customer orders. Use an AWS Lambda function to capture and store unprocessed orders.  
B. Use a Network Load Balancer (NLB) to manage web traffic. Use an Application Load Balancer to receive customer orders from the NLB. Use Amazon Redshift with a Multi-AZ deployment to store unprocessed and processed customer orders.  
C. Use a Gateway Load Balancer (GWLB) to manage web traffic. Use Amazon Elastic Container Service (Amazon ECS) to receive and process customer orders. Use the GWLB to capture and store unprocessed orders. Use Amazon DynamoDB to store processed customer orders.  
D. Use an Application Load Balancer to manage web traffic. Use Amazon EC2 Auto Scaling groups to receive and process customer orders. Use Amazon Simple Queue Service (Amazon SQS) to store unprocessed orders. Use Amazon RDS with a Multi-AZ deployment to store processed customer orders.

Answer: D  
Explanation: This architecture uses ALB for routing, Auto Scaling for elasticity, SQS to buffer unprocessed orders and decouple services, and RDS Multi-AZ for high availability and durability of transactional data. This ensures resilience and fault tolerance even during traffic spikes.

QUESTION 1017A company is using AWS DataSync to migrate millions of files from an on-premises system to AWS. The files are 10 KB in size on average. The company wants to use Amazon S3 for file storage. For the first year after the migration, the files will be accessed once or twice and must be immediately available. After 1 year, the files must be archived for at least 7 years. Which solution will meet these requirements MOST cost-effectively?  
A. Use an archive tool to group the files into large objects. Use DataSync to migrate the objects. Store the objects in S3 Glacier Instant Retrieval for the first year. Use a lifecycle configuration to transition the files to S3 Glacier Deep Archive after 1 year with a retention period of 7 years.  
B. Use an archive tool to group the files into large objects. Use DataSync to copy the objects to S3 Standard-Infrequent Access (S3 Standard-IA). Use a lifecycle configuration to transition the files to S3 Glacier Instant Retrieval after 1 year with a retention period of 7 years.

C. Configure the destination storage class for the files as S3 Glacier Instant Retrieval. Use a lifecycle policy to transition the files to S3 Glacier Flexible Retrieval after 1 year with a retention period of 7 years.  
D. Configure a DataSync task to transfer the files to S3 Standard-Infrequent Access (S3 Standard-IA). Use a lifecycle configuration to transition the files to S3 Deep Archive after 1 year with a retention period of 7 years.

Answer: D  
Explanation: QUESTION 1018A company recently performed a lift and shift migration of its on-premises Oracle database workload to run on an Amazon EC2 memory optimized Linux instance. The EC2 Linux instance uses a 1 TB Provisioned IOPS SSD (io1) EBS volume with 64,000 IOPS. The database storage performance after the migration is slower than the performance of the on-premises database. Which solution will improve storage performance?  
A. Add more Provisioned IOPS SSD (io1) EBS volumes. Use OS commands to create a Logical Volume Management (LVM) stripe.  
B. Increase

the Provisioned IOPS SSD (io1) EBS volume to more than 64,000 IOPS.C. Increase the size of the Provisioned IOPS SSD (io1) EBS volume to 2 TB.D. Change the EC2 Linux instance to a storage optimized instance type. Do not change the Provisioned IOPS SSD (io1) EBS volume.

**Answer:** AExplanation: The maximum provisioned IOPS for io1 is 64000 and hence you can achieve higher aggregate performance by adding more io1 volumes.

**QUESTION 1019** A company is migrating from a monolithic architecture for a web application that is hosted on Amazon EC2 to a serverless microservices architecture. The company wants to use AWS services that support an event-driven, loosely coupled architecture. The company wants to use the publish/subscribe (pub/sub) pattern. Which solution will meet these requirements MOST cost-effectively?

- A. Configure an Amazon API Gateway REST API to invoke an AWS Lambda function that publishes events to an Amazon Simple Queue Service (Amazon SQS) queue. Configure one or more subscribers to read events from the SQS queue.
- B. Configure an Amazon API Gateway REST API to invoke an AWS Lambda function that publishes events to an Amazon Simple Notification Service (Amazon SNS) topic. Configure one or more subscribers to receive events from the SNS topic.
- C. Configure an Amazon API Gateway WebSocket API to write to a data stream in Amazon Kinesis Data Streams with enhanced fan-out. Configure one or more subscribers to receive events from the data stream.
- D. Configure an Amazon API Gateway HTTP API to invoke an AWS Lambda function that publishes events to an Amazon Simple Notification Service (Amazon SNS) topic. Configure one or more subscribers to receive events from the topic.

**Answer:** B

**QUESTION 1020** A company recently migrated a monolithic application to an Amazon EC2 instance and Amazon RDS. The application has tightly coupled modules. The existing design of the application gives the application the ability to run on only a single EC2 instance. The company has noticed high CPU utilization on the EC2 instance during peak usage times. The high CPU utilization corresponds to degraded performance on Amazon RDS for read requests. The company wants to reduce the high CPU utilization and improve read request performance. Which solution will meet these requirements?

- A. Resize the EC2 instance to an EC2 instance type that has more CPU capacity. Configure an Auto Scaling group with a minimum and maximum size of 1.
- B. Configure an RDS read replica for read requests.
- C. Resize the EC2 instance to an EC2 instance type that has more CPU capacity. Configure an Auto Scaling group with a minimum and maximum size of 1. Add an RDS read replica and redirect all read/write traffic to the replica.
- D. Configure an Auto Scaling group with a minimum size of 1 and maximum size of 2. Resize the RDS DB instance to an instance type that has more CPU capacity.
- E. Resize the EC2 instance to an EC2 instance type that has more CPU capacity. Configure an Auto Scaling group with a minimum and maximum size of 1. Resize the RDS DB instance to an instance type that has more CPU capacity.

**Answer:** BExplanation: This approach addresses both the high CPU utilization on the EC2 instance and the degraded read performance on the RDS instance effectively.

**QUESTION 1021** A company needs to grant a team of developers access to the company's AWS resources. The company must maintain a high level of security for the resources. The company requires an access control solution that will prevent unauthorized access to the sensitive data. Which solution will meet these requirements?

- A. Share the IAM user credentials for each development team member with the rest of the team to simplify access management and to streamline development workflows.
- B. Define IAM roles that have fine-grained permissions based on the principle of least privilege. Assign an IAM role to each developer.
- C. Create IAM access keys to grant programmatic access to AWS resources. Allow only developers to interact with AWS resources through API calls by using the access keys.
- D. Create an AWS Cognito user pool. Grant developers access to AWS resources by using the user pool.

**Answer:** B

**QUESTION 1022** A company hosts a monolithic web application on an Amazon EC2 instance. Application users have recently reported poor performance at specific times. Analysis of Amazon CloudWatch metrics shows that CPU utilization is 100% during the periods of poor performance. The company wants to resolve this performance issue and improve application availability. Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

- A. Use AWS Compute Optimizer to obtain a recommendation for an instance type to scale vertically.
- B. Create an Amazon Machine Image (AMI) from the web server. Reference the AMI in a new launch template.
- C. Create an Auto Scaling group and an Application Load Balancer to scale vertically.
- D. Use AWS Compute Optimizer to obtain a recommendation for an instance type to scale horizontally.
- E. Create an Auto Scaling group and an Application Load Balancer to scale horizontally.

**Answer:** BE

**QUESTION 1023** A company runs all its business applications in the AWS Cloud. The company uses AWS Organizations to manage multiple AWS accounts. A solutions architect needs to review all permissions that are granted to IAM users to determine which IAM users have more permissions than required. Which solution will meet these requirements with the LEAST administrative overhead?

- A. Use Network Access Analyzer to review all access permissions in the company's AWS accounts.
- B. Create an AWS CloudWatch alarm that activates when an IAM user creates or modifies resources in an AWS account.
- C. Use AWS Identity and Access Management (IAM) Access Analyzer to review all the company's resources and accounts.
- D. Use Amazon Inspector to find vulnerabilities in existing IAM policies.

**Answer:** C

**QUESTION 1024** A company needs to implement a new data retention policy for regulatory compliance. As part of this policy, sensitive documents that are stored in an Amazon S3 bucket must be protected from deletion or modification for a fixed period of

time.Which solution will meet these requirements?A. Activate S3 Object Lock on the required objects and enable governance mode.B. Activate S3 Object Lock on the required objects and enable compliance mode.C. Enable versioning on the S3 bucket. Set a lifecycle policy to delete the objects after a specified period.D. Configure an S3 Lifecycle policy to transition objects to S3 Glacier Flexible Retrieval for the retention duration.

Answer: B

QUESTION 1025A company runs its customer-facing web application on containers. The workload uses Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. The web application is resource intensive. The web application needs to be available 24 hours a day, 7 days a week for customers. The company expects the application to experience short bursts of high traffic. The workload must be highly available. Which solution will meet these requirements MOST cost-effectively?A. Configure an ECS capacity provider with Fargate. Conduct load testing by using a third-party tool. Rightsize the Fargate tasks in Amazon CloudWatch.B. Configure an ECS capacity provider with Fargate for steady state and Fargate Spot for burst traffic.C. Configure an ECS capacity provider with Fargate Spot for steady state and Fargate for burst traffic.D. Configure an ECS capacity provider with Fargate. Use AWS Compute Optimizer to rightsize the Fargate task.

Answer: B

Explanation: This combination leverages the cost benefits of Fargate Spot for burst traffic while ensuring steady performance with regular Fargate instances.

QUESTION 1026A company is building an application in the AWS Cloud. The application is hosted on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses Amazon Route 53 for the DNS. The company needs a managed solution with proactive engagement to detect against DDoS attacks. Which solution will meet these requirements?A. Enable AWS Config. Configure an AWS Config managed rule that detects DDoS attacks.B. Enable AWS WAF on the ALB. Create an AWS WAF web ACL with rules to detect and prevent DDoS attacks. Associate the web ACL with the ALB.C. Store the ALB access logs in an Amazon S3 bucket. Configure Amazon GuardDuty to detect and take automated preventative actions for DDoS attacks.D. Subscribe to AWS Shield Advanced. Configure hosted zones in Route 53. Add ALB resources as protected resources.

Answer: D

QUESTION 1027A company hosts a video streaming web application in a VPC. The company uses a Network Load Balancer (NLB) to handle TCP traffic for real-time data processing. There have been unauthorized attempts to access the application. The company wants to improve application security with minimal architectural change to prevent unauthorized attempts to access the application. Which solution will meet these requirements?A. Implement a series of AWS WAF rules directly on the NLB to filter out unauthorized traffic.B. Recreate the NLB with a security group to allow only trusted IP addresses.C. Deploy a second NLB in parallel with the existing NLB configured with a strict IP address allow list.D. Use AWS Shield Advanced to provide enhanced DDoS protection and prevent unauthorized access attempts.

Answer: B

QUESTION 1028A healthcare company is developing an AWS Lambda function that publishes notifications to an encrypted Amazon Simple Notification Service (Amazon SNS) topic. The notifications contain protected health information (PHI). The SNS topic uses AWS Key Management Service (AWS KMS) customer managed keys for encryption. The company must ensure that the application has the necessary permissions to publish messages securely to the SNS topic. Which combination of steps will meet these requirements? (Choose three.)A. Create a resource policy for the SNS topic that allows the Lambda function to publish messages to the topic.B. Use server-side encryption with AWS KMS keys (SSE-KMS) for the SNS topic instead of customer managed keys.C. Create a resource policy for the encryption key that the SNS topic uses that has the necessary AWS KMS permissions.D. Specify the Lambda function's Amazon Resource Name (ARN) in the SNS topic's resource policy.E. Associate an Amazon API Gateway HTTP API with the SNS topic to control access to the topic by using API Gateway resource policies.F. Configure a Lambda execution role that has the necessary IAM permissions to use a customer managed key in AWS KMS.

Answer: ACF

QUESTION 1029A company has an employee web portal. Employees log in to the portal to view payroll details. The company is developing a new system to give employees the ability to upload scanned documents for reimbursement. The company runs a program to extract text-based data from the documents and attach the extracted information to each employee's reimbursement IDs for processing. The employee web portal requires 100% uptime. The document extract program runs infrequently throughout the day on an on-demand basis. The company wants to build a scalable and cost-effective new system that will require minimal changes to the existing web portal. The company does not want to make any code changes. Which solution will meet these requirements with the LEAST implementation effort?A. Run Amazon EC2 On-Demand Instances in an Auto Scaling group for the web portal. Use an AWS Lambda function to run the document extract program. Invoke the Lambda function when an employee uploads a new reimbursement document.B. Run Amazon EC2 Spot Instances in an Auto Scaling group for the web portal. Run the document extract program on EC2 Spot Instances. Start document extract program instances when an employee uploads a new reimbursement document.C. Purchase a Savings Plan to run the web portal and the document extract program. Run the web portal and the document extract program in an Auto Scaling group.D. Create an Amazon S3 bucket to host the web portal. Use Amazon API Gateway and an AWS Lambda function for the existing functionalities. Use the Lambda function to run the document extract program. Invoke the Lambda function when the API that is associated with a new document upload is called.

Answer: A

1030A media company has a multi-account AWS environment in the us-east-1 Region. The company has an Amazon Simple Notification Service (Amazon SNS) topic in a production account that publishes performance metrics. The company has an AWS Lambda function in an administrator account to process and analyze log data. The Lambda function that is in the administrator account must be invoked by messages from the SNS topic that is in the production account when significant metrics are reported. Which combination of steps will meet these requirements? (Choose two.)

- A. Create an IAM resource policy for the Lambda function that allows Amazon SNS to invoke the function.
- B. Implement an Amazon Simple Queue Service (Amazon SQS) queue in the administrator account to buffer messages from the SNS topic that is in the production account. Configure the SQS queue to invoke the Lambda function.
- C. Create an IAM policy for the SNS topic that allows the Lambda function to subscribe to the topic.
- D. Use an Amazon EventBridge rule in the production account to capture the SNS topic notifications. Configure the EventBridge rule to forward notifications to the Lambda function that is in the administrator account.
- E. Store performance metrics in an Amazon S3 bucket in the production account. Use Amazon Athena to analyze the metrics from the administrator account.

Answer: AC

QUESTION 1031A company is migrating an application from an on-premises location to Amazon Elastic Kubernetes Service (Amazon EKS). The company must use a custom subnet for pods that are in the company's VPC to comply with requirements. The company also needs to ensure that the pods can communicate securely within the pods' VPC. Which solution will meet these requirements?

- A. Configure AWS Transit Gateway to directly manage custom subnet configurations for the pods in Amazon EKS.
- B. Create an AWS Direct Connect connection from the company's on-premises IP address ranges to the EKS pods.
- C. Use the Amazon VPC CNI plugin for Kubernetes. Define custom subnets in the VPC cluster for the pods to use.
- D. Implement a Kubernetes network policy that has pod anti-affinity rules to restrict pod placement to specific nodes that are within custom subnets.

Answer: C

Explanation: The Amazon VPC Container Network Interface (CNI) plugin is the default network plugin for Amazon EKS. It allows Kubernetes pods to receive IP addresses from a VPC's subnet and enables pods to communicate securely within the VPC as if they were native VPC resources.

QUESTION 1032A company hosts an ecommerce application that stores all data in a single Amazon RDS for MySQL DB instance that is fully managed by AWS. The company needs to mitigate the risk of a single point of failure. Which solution will meet these requirements with the LEAST implementation effort?

- A. Modify the RDS DB instance to use a Multi-AZ deployment. Apply the changes during the next maintenance window.
- B. Migrate the current database to a new Amazon DynamoDB Multi-AZ deployment. Use AWS Database Migration Service (AWS DMS) with a heterogeneous migration strategy to migrate the current RDS DB instance to DynamoDB tables.
- C. Create a new RDS DB instance in a Multi-AZ deployment. Manually restore the data from the existing RDS DB instance from the most recent snapshot.
- D. Configure the DB instance in an Amazon EC2 Auto Scaling group with a minimum group size of three. Use Amazon Route 53 simple routing to distribute requests to all DB instances.

Answer: A

QUESTION 1033A company has multiple Microsoft Windows SMB file servers and Linux NFS file servers for file sharing in an on-premises environment. As part of the company's AWS migration plan, the company wants to consolidate the file servers in the AWS Cloud. The company needs a managed AWS storage service that supports both NFS and SMB access. The solution must be able to share between protocols. The solution must have redundancy at the Availability Zone level. Which solution will meet these requirements?

- A. Use Amazon FSx for NetApp ONTAP for storage. Configure multi-protocol access.
- B. Create two Amazon EC2 instances. Use one EC2 instance for Windows SMB file server access and one EC2 instance for Linux NFS file server access.
- C. Use Amazon FSx for NetApp ONTAP for SMB access. Use Amazon FSx for Lustre for NFS access.
- D. Use Amazon S3 storage. Access Amazon S3 through an Amazon S3 File Gateway.

Answer: A

QUESTION 1034A software company needs to upgrade a critical web application. The application currently runs on a single Amazon EC2 instance that the company hosts in a public subnet. The EC2 instance runs a MySQL database. The application's DNS records are published in an Amazon Route 53 zone. A solutions architect must reconfigure the application to be scalable and highly available. The solutions architect must also reduce MySQL read latency. Which combination of solutions will meet these requirements? (Choose two.)

- A. Launch a second EC2 instance in a second AWS Region. Use a Route 53 failover routing policy to redirect the traffic to the second EC2 instance.
- B. Create and configure an Auto Scaling group to launch private EC2 instances in multiple Availability Zones. Add the instances to a target group behind a new Application Load Balancer.
- C. Migrate the database to an Amazon Aurora MySQL cluster. Create the primary DB instance and reader DB instance in separate Availability Zones.
- D. Create and configure an Auto Scaling group to launch private EC2 instances in multiple AWS Regions. Add the instances to a target group behind a new Application Load Balancer.
- E. Migrate the database to an Amazon Aurora MySQL cluster with cross-Region read replicas.

Answer: B

Explanation: To improve scalability and availability, EC2 Auto Scaling across multiple Availability Zones with an Application Load Balancer ensures resilient infrastructure. Migrating to Amazon Aurora MySQL with reader endpoints reduces read latency by offloading read traffic to replicas in other AZs, while also increasing high availability.

QUESTION 1035A company runs thousands of AWS Lambda functions. The company needs a solution to securely

store sensitive information that all the Lambda functions use. The solution must also manage the automatic rotation of the sensitive information. Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose two.)

- A. Create HTTP security headers by using Lambda@Edge to retrieve and create sensitive information
- B. Create a Lambda layer that retrieves sensitive information
- C. Store sensitive information in AWS Secrets Manager
- D. Store sensitive information in AWS Systems Manager Parameter Store
- E. Create a Lambda consumer with dedicated throughput to retrieve sensitive information and create environmental variables

Answer: BCEExplanation: AWS Secrets Manager securely stores sensitive information and provides automatic rotation of secrets, reducing the need for manual management. Using a Lambda layer allows multiple Lambda functions to access the sensitive information stored in Secrets Manager without needing to duplicate retrieval logic in each function. This approach centralizes the retrieval process and reduces operational complexity.

QUESTION 1036

A company has an internal application that runs on Amazon EC2 instances in an Auto Scaling group. The EC2 instances are compute optimized and use Amazon Elastic Block Store (Amazon EBS) volumes. The company wants to identify cost optimizations across the EC2 instances, the Auto Scaling group, and the EBS volumes. Which solution will meet these requirements with the MOST operational efficiency?

- A. Create a new AWS Cost and Usage Report. Search the report for cost recommendations for the EC2 instances, the Auto Scaling group, and the EBS volumes.
- B. Create new Amazon CloudWatch billing alerts. Check the alert statuses for cost recommendations for the EC2 instances, the Auto Scaling group, and the EBS volumes.
- C. Configure AWS Compute Optimizer for cost recommendations for the EC2 instances, the Auto Scaling group and the EBS volumes.
- D. Configure AWS Compute Optimizer for cost recommendations for the EC2 instances. Create a new AWS Cost and Usage Report. Search the report for cost recommendations for the Auto Scaling group and the EBS volumes.

Answer: C

QUESTION 1037

A company is running a media store across multiple Amazon EC2 instances distributed across multiple Availability Zones in a single VPC. The company wants a high-performing solution to share data between all the EC2 instances, and prefers to keep the data within the VPC only. What should a solutions architect recommend?

- A. Create an Amazon S3 bucket and call the service APIs from each instance's application
- B. Create an Amazon S3 bucket and configure all instances to access it as a mounted volume
- C. Configure an Amazon Elastic Block Store (Amazon EBS) volume and mount it across all instances
- D. Configure an Amazon Elastic File System (Amazon EFS) file system and mount it across all instances

Answer: D

QUESTION 1038

A company uses an Amazon RDS for MySQL instance. To prepare for end-of-year processing, the company added a read replica to accommodate extra read-only queries from the company's reporting tool. The read replica CPU usage was 60% and the primary instance CPU usage was 60%. After end-of-year activities are complete, the read replica has a constant 25% CPU usage. The primary instance still has a constant 60% CPU usage. The company wants to rightsize the database and still provide enough performance for future growth. Which solution will meet these requirements?

- A. Delete the read replica
- B. Do not make changes to the primary instance
- C. Resize the read replica to a smaller instance size
- D. Do not make changes to the primary instance
- E. Resize the read replica to a larger instance size
- F. Resize the primary instance to a smaller instance size
- G. Delete the read replica
- H. Resize the primary instance to a larger instance

Answer: B

QUESTION 1039

A company is migrating its databases to Amazon RDS for PostgreSQL. The company is migrating its applications to Amazon EC2 instances. The company wants to optimize costs for long-running workloads. Which solution will meet this requirement MOST cost-effectively?

- A. Use On-Demand Instances for the Amazon RDS for PostgreSQL workloads. Purchase a 1 year Compute Savings Plan with the No Upfront option for the EC2 instances.
- B. Purchase Reserved Instances for a 1 year term with the No Upfront option for the Amazon RDS for PostgreSQL workloads. Purchase a 1 year EC2 Instance Savings Plan with the No Upfront option for the EC2 instances.
- C. Purchase Reserved Instances for a 1 year term with the Partial Upfront option for the Amazon RDS for PostgreSQL workloads. Purchase a 1 year EC2 Instance Savings Plan with the Partial Upfront option for the EC2 instances.
- D. Purchase Reserved Instances for a 3 year term with the All Upfront option for the Amazon RDS for PostgreSQL workloads. Purchase a 3 year EC2 Instance Savings Plan with the All Upfront option for the EC2 instances.

Answer: D

QUESTION 1040

A company is using an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The company must ensure that Kubernetes service accounts in the EKS cluster have secure and granular access to specific AWS resources by using IAM roles for service accounts (IRSA). Which combination of solutions will meet these requirements? (Choose two.)

- A. Create an IAM policy that defines the required permissions
- B. Attach the policy directly to the IAM role of the EKS nodes
- C. Implement network policies within the EKS cluster to prevent Kubernetes service accounts from accessing specific AWS services
- D. Modify the EKS cluster's IAM role to include permissions for each Kubernetes service account
- E. Ensure a one-to-one mapping between IAM roles and Kubernetes roles
- F. Define an IAM role that includes the necessary permissions
- G. Annotate the Kubernetes service accounts with the Amazon ResourceName (ARN) of the IAM role
- H. Set up a trust relationship between the IAM roles for the service accounts and an OpenID Connect (OIDC) identity provider

Answer: DE

QUESTION 1041

A company regularly uploads confidential data to Amazon S3 buckets for analysis. The company's security policies mandate that the objects must be encrypted at rest. The company must

automatically rotate the encryption key every year. The company must be able to track key rotation by using AWS CloudTrail. The company also must minimize costs for the encryption key. Which solution will meet these requirements?

A. Use server-side encryption with customer-provided keys (SSE-C)

B. Use server-side encryption with Amazon S3 managed keys (SSE-S3)

C. Use server-side encryption with AWS KMS keys (SSE-KMS)

D. Use server-side encryption with customer managed AWS KMS keys

Answer: C

Explanation: <https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html#aws-managed-cmk>

QUESTION 1042

A company has migrated several applications to AWS in the past 3 months. The company wants to know the breakdown of costs for each of these applications. The company wants to receive a regular report that includes this information. Which solution will meet these requirements MOST cost-effectively?

A. Use AWS Budgets to download data for the past 3 months into a .csv file. Look up the desired information.

B. Load AWS Cost and Usage Reports into an Amazon RDS DB instance. Run SQL queries to get the desired information.

C. Tag all the AWS resources with a key for cost and a value of the application's name. Activate cost allocation tags. Use Cost Explorer to get the desired information.

D. Tag all the AWS resources with a key for cost and a value of the application's name. Use the AWS Billing and Cost Management console to download bills for the past 3 months. Look up the desired information.

Answer: C

QUESTION 1043

An ecommerce company is preparing to deploy a web application on AWS to ensure continuous service for customers. The architecture includes a web application that the company hosts on Amazon EC2 instances, a relational database in Amazon RDS, and static assets that the company stores in Amazon S3. The company wants to design a robust and resilient architecture for the application. Which solution will meet these requirements?

A. Deploy Amazon EC2 instances in a single Availability Zone. Deploy an RDS DB instance in the same Availability Zone. Use Amazon S3 with versioning enabled to store static assets.

B. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Deploy a Multi-AZ RDS DB instance. Use Amazon CloudFront to distribute static assets.

C. Deploy Amazon EC2 instances in a single Availability Zone. Deploy an RDS DB instance in a second Availability Zone for cross-AZ redundancy. Serve static assets directly from the EC2 instances.

D. Use AWS Lambda functions to serve the web application. Use Amazon Aurora Serverless v2 for the database. Store static assets in Amazon Elastic File System (Amazon EFS) One Zone-Infrequent Access (One Zone-IA).

Answer: B

QUESTION 1044

An ecommerce company runs several internal applications in multiple AWS accounts. The company uses AWS Organizations to manage its AWS accounts. A security appliance in the company's networking account must inspect interactions between applications across AWS accounts. Which solution will meet these requirements?

A. Deploy a Network Load Balancer (NLB) in the networking account to send traffic to the security appliance. Configure the application accounts to send traffic to the NLB by using an interface VPC endpoint in the application accounts.

B. Deploy an Application Load Balancer (ALB) in the application accounts to send traffic directly to the security appliance.

C. Deploy a Gateway Load Balancer (GWLB) in the networking account to send traffic to the security appliance. Configure the application accounts to send traffic to the GWLB by using an interface GWLB endpoint in the application accounts.

D. Deploy an interface VPC endpoint in the application accounts to send traffic directly to the security appliance.

Answer: C

QUESTION 1045

A company runs its production workload on an Amazon Aurora MySQL DB cluster that includes six Aurora Replicas. The company wants near-real-time reporting queries from one of its departments to be automatically distributed across three of the Aurora Replicas. Those three replicas have a different compute and memory specification from the rest of the DB cluster. Which solution meets these requirements?

A. Create and use a custom endpoint for the workload.

B. Create a three-node cluster clone and use the reader endpoint.

C. Use any of the instance endpoints for the selected three nodes.

D. Use the reader endpoint to automatically distribute the read-only workload.

Answer: B

QUESTION 1046

A company runs a Node.js function on a server in its on-premises data center. The data center stores data in a PostgreSQL database. The company stores the credentials in a connection string in an environment variable on the server. The company wants to migrate its application to AWS and to replace the Node.js application server with AWS Lambda. The company also wants to migrate to Amazon RDS for PostgreSQL and to ensure that the database credentials are securely managed. Which solution will meet these requirements with the LEAST operational overhead?

A. Store the database credentials as a parameter in AWS Systems Manager Parameter Store. Configure Parameter Store to automatically rotate the secrets every 30 days. Update the Lambda function to retrieve the credentials from the parameter.

B. Store the database credentials as a secret in AWS Secrets Manager. Configure Secrets Manager to automatically rotate the credentials every 30 days. Update the Lambda function to retrieve the credentials from the secret.

C. Store the database credentials as an encrypted Lambda environment variable. Write a custom Lambda function to rotate the credentials. Schedule the Lambda function to run every 30 days.

D. Store the database credentials as a key in AWS Key Management Service (AWS KMS). Configure automatic rotation for the key. Update the Lambda function to retrieve the credentials from the KMS key.

Answer: B

QUESTION 1047

A company wants to replicate existing and ongoing data changes from an on-premises Oracle database to Amazon RDS for Oracle. The amount of data to replicate varies throughout each day. The company wants to use AWS Database Migration Service (AWS DMS) for data replication. The solution must allocate only

the capacity that the replication instance requires. Which solution will meet these requirements?

A. Configure the AWS DMS replication instance with a Multi-AZ deployment to provision instances across multiple Availability Zones.

B. Create an AWS DMS Serverless replication task to analyze and replicate the data while provisioning the required capacity.

C. Use Amazon EC2 Auto Scaling to scale the size of the AWS DMS replication instance up or down based on the amount of data to replicate.

D. Provision AWS DMS replication capacity by using Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch type to analyze and replicate the data while provisioning the required capacity.

Answer: BExplanation: AWS DMS Serverless is designed to automatically allocate and manage the necessary compute and memory resources based on the demand of the data replication workload. It scales capacity up or down according to the data replication requirements without manual intervention. This approach ensures that the replication task uses only the required capacity at any given time, optimizing costs and resources, especially given that the amount of data to replicate varies throughout the day.

QUESTION 1048

A company has a multi-tier web application. The application's internal service components are deployed on Amazon EC2 instances. The internal service components need to access third-party software as a service (SaaS) APIs that are hosted on AWS. The company needs to provide secure and private connectivity from the application's internal services to the third-party SaaS application. The company needs to ensure that there is minimal public internet exposure. Which solution will meet these requirements?

A. Implement an AWS Site-to-Site VPN to establish a secure connection with the third-party SaaS provider.

B. Deploy AWS Transit Gateway to manage and route traffic between the application's VPC and the third-party SaaS provider.

C. Configure AWS PrivateLink to allow only outbound traffic from the VPC without enabling the third-party SaaS provider to establish.

D. Use AWS PrivateLink to create a private connection between the application's VPC and the third-party SaaS provider.

Answer: DExplanation: AWS PrivateLink enables private connectivity between VPCs and supported AWS or third-party services without exposing traffic to the public internet. It ensures secure and private communications, making it ideal for connecting internal services to SaaS applications hosted in AWS.

QUESTION 1049

A solutions architect needs to connect a company's corporate network to its VPC to allow on-premises access to its AWS resources. The solution must provide encryption of all traffic between the corporate network and the VPC at the network layer and the session layer. The solution also must provide security controls to prevent unrestricted access between AWS and the on-premises systems. Which solution meets these requirements?

A. Configure AWS Direct Connect to connect to the VPC. Configure the VPC route tables to allow and deny traffic between AWS and on-premises as required.

B. Create an IAM policy to allow access to the AWS Management Console only from a defined set of corporate IP addresses. Restrict user access based on job responsibility by using an IAM policy and roles.

C. Configure AWS Site-to-Site VPN to connect to the VPC. Configure route table entries to direct traffic from on-premises to the VPC.

D. Configure AWS Transit Gateway to connect to the VPC. Configure route table entries to direct traffic from on-premises to the VPC.

Answer: CExplanation: AWS Direct Connect does not provide encryption by itself; it is often used in conjunction with VPN for encrypted traffic. Direct Connect primarily offers a dedicated connection and does not inherently satisfy the encryption requirement.

QUESTION 1050

A company has a custom application with embedded credentials that retrieves information from a database in an Amazon RDS for MySQL DB cluster. The company needs to make the application more secure with minimal programming effort. The company has created credentials on the RDS for MySQL database for the application user. Which solution will meet these requirements?

A. Store the credentials in AWS Key Management Service (AWS KMS). Create keys in AWS KMS. Configure the application to load the database credentials from AWS KMS. Enable automatic key rotation.

B. Store the credentials in encrypted local storage. Configure the application to load the database credentials from the local storage. Set up a credentials rotation schedule by creating a cron job.

C. Store the credentials in AWS Secrets Manager. Configure the application to load the database credentials from Secrets Manager. Set up a credentials rotation schedule by creating an AWS Lambda function for Secrets Manager.

D. Store the credentials in AWS Systems Manager Parameter Store. Configure the application to load the database credentials from Parameter Store. Set up a credentials rotation schedule in the RDS for MySQL database by using Parameter Store.

Answer: C

QUESTION 1051

A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing data and new data by using SQL. The company stores the data in an Amazon S3 bucket. The data must be encrypted at rest and replicated to a different AWS Region. Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new S3 bucket that uses server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Configure Cross-Region Replication (CRR). Load the data into the new S3 bucket. Use Amazon Athena to query the data.

B. Create a new S3 bucket that uses server-side encryption with Amazon S3 managed keys (SSE-S3). Configure Cross-Region Replication (CRR). Load the data into the new S3 bucket. Use Amazon RDS to query the data.

C. Configure Cross-Region Replication (CRR) on the existing S3 bucket. Use server-side encryption with Amazon S3 managed keys (SSE-S3). Use Amazon

Athena to query the data.D. Configure S3 Cross-Region Replication (CRR) on the existing S3 bucket. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.Answer: AQUESTION 1052A company has a web application that has thousands of users. The application uses 8-10 user-uploaded images to generate AI images. Users can download the generated AI images once every 6 hours. The company also has a premium user option that gives users the ability to download the generated AI images anytime.The company uses the user-uploaded images to run AI model training twice a year. The company needs a storage solution to store the images.Which storage solution meets these requirements MOST cost-effectively?A. Move uploaded images to Amazon S3 Glacier Deep Archive. Move premium user-generated AI images to S3 Standard. Move non-premium user-generated AI images to S3 Standard-Infrequent Access (S3 Standard-IA).B. Move uploaded images to Amazon S3 Glacier Deep Archive Move all generated AI images to S3 Glacier Flexible Retrieval.C. Move uploaded images to Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA). Move premium user-generated AI images to S3 Standard. Move non-premium user-generated AI images to S3 Standard-Infrequent Access (S3 Standard-IA).D. Move uploaded images to Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA). Move all generated AI images to S3 Glacier Flexible Retrieval.Answer: AQUESTION 1053A company is developing machine learning (ML) models on AWS. The company is developing the ML models as independent microservices. The microservices fetch approximately 1 GB of model data from Amazon S3 at startup and load the data into memory. Users access the ML models through an asynchronous API. Users can send a request or a batch of requests.The company provides the ML models to hundreds of users. The usage patterns for the models are irregular. Some models are not used for days or weeks. Other models receive batches of thousands of requests at a time.Which solution will meet these requirements?A. Direct the requests from the API to a Network Load Balancer (NLB). Deploy the ML models as AWS Lambda functions that the NLB will invoke. Use auto scaling to scale the Lambda functions based on the traffic that the NLB receives.B. Direct the requests from the API to an Application Load Balancer (ALB). Deploy the ML models as Amazon Elastic Container Service (Amazon ECS) services that the ALB will invoke. Use auto scaling to scale the ECS cluster instances based on the traffic that the ALB receives.C. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue. Deploy the ML models as AWS Lambda functions that SQS events will invoke. Use auto scaling to increase the number of vCPUs for the Lambda functions based on the size of the SQS queue.D. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue. Deploy the ML models as Amazon Elastic Container Service (Amazon ECS) services that read from the queue. Use auto scaling for Amazon ECS to scale both the cluster capacity and number of the services based on the size of the SQS queue.Answer: DQUESTION 1054A company runs a web application on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB). The application stores data in an Amazon Aurora MySQL DB cluster.The company needs to create a disaster recovery (DR) solution. The acceptable recovery time for the DR solution is up to 30 minutes. The DR solution does not need to support customer usage when the primary infrastructure is healthy.Which solution will meet these requirements?A. Deploy the DR infrastructure in a second AWS Region with an ALB and an Auto Scaling group. Set the desired capacity and maximum capacity of the Auto Scaling group to a minimum value. Convert the Aurora MySQL DB cluster to an Aurora global database. Configure Amazon Route 53 for an active-passive failover with ALB endpoints.B. Deploy the DR infrastructure in a second AWS Region with an ALB. Update the Auto Scaling group to include EC2 instances from the second Region. Use Amazon Route 53 to configure active-active failover. Convert the Aurora MySQL DB cluster to an Aurora global database.C. Back up the Aurora MySQL DB cluster data by using AWS Backup. Deploy the DR infrastructure in a second AWS Region with an ALB. Update the Auto Scaling group to include EC2 instances from the second Region. Use Amazon Route 53 to configure active-active failover. Create an Aurora MySQL DB cluster in the second Region Restore the data from the backup.D. Back up the infrastructure configuration by using AWS Backup. Use the backup to create the required infrastructure in a second AWS Region. Set the Auto Scaling group desired capacity to zero. Use Amazon Route 53 to configure active-passive failover. Convert the Aurora MySQL DB cluster to an Aurora global database.Answer: AQUESTION 1055A company is migrating its data processing application to the AWS Cloud. The application processes several short-lived batch jobs that cannot be disrupted. Data is generated after each batch job is completed. The data is accessed for 30 days and retained for 2 years.The company wants to keep the cost of running the application in the AWS Cloud as low as possible.Which solution will meet these requirements?A. Migrate the data processing application to Amazon EC2 Spot Instances. Store the data in Amazon S3 Standard. Move the data to Amazon S3 Glacier Instant. Retrieval after 30 days. Set an expiration to delete the data after 2 years.B. Migrate the data processing application to Amazon EC2 On-Demand Instances. Store the data in Amazon S3 Glacier Instant Retrieval. Move the data to S3 Glacier Deep Archive after 30 days. Set an expiration to delete the data after 2 years.C. Deploy Amazon EC2 Spot Instances to run the batch jobs. Store the data in Amazon S3 Standard. Move the data to Amazon S3 Glacier Flexible Retrieval after 30 days. Set an expiration to delete the data after 2 years.D. Deploy Amazon EC2 On-Demand Instances to run the batch jobs. Store the data in Amazon S3 Standard. Move the data to Amazon S3

Glacier Deep Archive after 30 days. Set an expiration to delete the data after 2 years. Answer: DQUESTION 1056A company needs to design a hybrid network architecture. The company's workloads are currently stored in the AWS Cloud and in on-premises data centers. The workloads require single-digit latencies to communicate. The company uses an AWS Transit Gateway transit gateway to connect multiple VPCs. Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.) A. Establish an AWS Site-to-Site VPN connection to each VPC. B. Associate an AWS Direct Connect gateway with the transit gateway that is attached to the VPCs. C. Establish an AWS Site-to-Site VPN connection to an AWS Direct Connect gateway. D. Establish an AWS Direct Connect connection. Create a transit virtual interface (VIF) to a Direct Connect gateway. E. Associate AWS Site-to-Site VPN connections with the transit gateway that is attached to the VPCs. Answer: BDQUESTION 1057A global ecommerce company runs its critical workloads on AWS. The workloads use an Amazon RDS for PostgreSQL DB instance that is configured for a Multi-AZ deployment. Customers have reported application timeouts when the company undergoes database failovers. The company needs a resilient solution to reduce failover time. Which solution will meet these requirements? A. Create an Amazon RDS Proxy. Assign the proxy to the DB instance. B. Create a read replica for the DB instance. Move the read traffic to the read replica. C. Enable Performance Insights. Monitor the CPU load to identify the timeouts. D. Take regular automatic snapshots. Copy the automatic snapshots to multiple AWS Regions. Answer: AQUESTION 1058A company has multiple Amazon RDS DB instances that run in a development AWS account. All the instances have tags to identify them as development resources. The company needs the development DB instances to run on a schedule only during business hours. Which solution will meet these requirements with the LEAST operational overhead? A. Create an Amazon CloudWatch alarm to identify RDS instances that need to be stopped. Create an AWS Lambda function to start and stop the RDS instances. B. Create an AWS Trusted Advisor report to identify RDS instances to be started and stopped. Create an AWS Lambda function to start and stop the RDS instances. C. Create AWS Systems Manager State Manager associations to start and stop the RDS instances. D. Create an Amazon EventBridge rule that invokes AWS Lambda functions to start and stop the RDS instances. Answer: CExplanation: AWS Systems Manager State Manager allows you to automate the process of starting and stopping RDS instances based on a defined schedule. QUESTION 1059 A consumer survey company has gathered data for several years from a specific geographic region. The company stores this data in an Amazon S3 bucket in an AWS Region. The company has started to share this data with a marketing firm in a new geographic region. The company has granted the firm's AWS account access to the S3 bucket. The company wants to minimize the data transfer costs when the marketing firm requests data from the S3 bucket. Which solution will meet these requirements? A. Configure the Requester Pays feature on the company's S3 bucket. B. Configure S3 Cross-Region Replication (CRR) from the company's S3 bucket to one of the marketing firm's S3 buckets. C. Configure AWS Resource Access Manager to share the S3 bucket with the marketing firm AWS account. D. Configure the company's S3 bucket to use S3 Intelligent-Tiering Sync the S3 bucket to one of the marketing firm's S3 buckets. Answer: AExplanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/RequesterPaysBuckets.html> QUESTION 1060A company uses AWS to host its public ecommerce website. The website uses an AWS Global Accelerator accelerator for traffic from the internet. The Global Accelerator accelerator forwards the traffic to an Application Load Balancer (ALB) that is the entry point for an Auto Scaling group. The company recently identified a DDoS attack on the website. The company needs a solution to mitigate future attacks. Which solution will meet these requirements with the LEAST implementation effort? A. Configure an AWS WAF web ACL for the Global Accelerator accelerator to block traffic by using rate-based rules. B. Configure an AWS Lambda function to read the ALB metrics to block attacks by updating a VPC network ACL. C. Configure an AWS WAF web ACL on the ALB to block traffic by using rate-based rules. D. Configure an Amazon CloudFront distribution in front of the Global Accelerator accelerator. Answer: CExplanation: <https://repost.aws/knowledge-center/globalaccelerator-aws-waf-filter-layer7-traffic> QUESTION 1061A company uses an Amazon DynamoDB table to store data that the company receives from devices. The DynamoDB table supports a customer-facing website to display recent activity on customer devices. The company configured the table with provisioned throughput for writes and reads. The company wants to calculate performance metrics for customer device data on a daily basis. The solution must have minimal effect on the table's provisioned read and write capacity. Which solution will meet these requirements? A. Use an Amazon Athena SQL query with the Amazon Athena DynamoDB connector to calculate performance metrics on a recurring schedule. B. Use an AWS Glue job with the AWS Glue DynamoDB export connector to calculate performance metrics on a recurring schedule. C. Use an Amazon Redshift COPY command to calculate performance metrics on a recurring schedule. D. Use an Amazon EMR job with an Apache Hive external table to calculate performance metrics on a recurring schedule. Answer: BExplanation: DynamoDB export connector literally "exports" table snapshot to s3 as dynamoDB-json object, then process on it. So it does not affect on read / write capacity on dynamoDB itself. But Athena query directly on dynamoDB so affects on read / write capacity. QUESTION 1062A solutions architect is designing the cloud architecture for a new stateless

application that will be deployed on AWS. The solutions architect created an Amazon Machine Image (AMI) and launch template for the application. Based on the number of jobs that need to be processed, the processing must run in parallel while adding and removing application Amazon EC2 instances as needed. The application must be loosely coupled. The job items must be durably stored. Which solution will meet these requirements?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic to send the jobs that need to be processed. Create an Auto Scaling group by using the launch template with the scaling policy set to add and remove EC2 instances based on CPU usage.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue to hold the jobs that need to be processed. Create an Auto Scaling group by using the launch template with the scaling policy set to add and remove EC2 instances based on network usage.
- C. Create an Amazon Simple Queue Service (Amazon SQS) queue to hold the jobs that need to be processed. Create an Auto Scaling group by using the launch template with the scaling policy set to add and remove EC2 instances based on the number of items in the SQS queue.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic to send the jobs that need to be processed. Create an Auto Scaling group by using the launch template with the scaling policy set to add and remove EC2 instances based on the number of messages published to the SNS topic.

Answer: C

Explanation: Amazon SQS provides durable, decoupled message storage for distributed systems. Using SQS as a job queue enables each EC2 instance to process a message independently. Scaling the Auto Scaling group based on the SQS queue length ensures parallelism and elasticity, aligning compute resources with workload volume.

**QUESTION 1063** A global ecommerce company uses a monolithic architecture. The company needs a solution to manage the increasing volume of product data. The solution must be scalable and have a modular service architecture. The company needs to maintain its structured database schemas. The company also needs a storage solution to store product data and product images. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use an Amazon EC2 instance in an Auto Scaling group to deploy a containerized application. Use an Application Load Balancer to distribute web traffic. Use an Amazon RDS DB instance to store product data and product images.
- B. Use AWS Lambda functions to manage the existing monolithic application. Use Amazon DynamoDB to store product data and product images. Use Amazon Simple Notification Service (Amazon SNS) for event-driven communication between the Lambda functions.
- C. Use Amazon Elastic Kubernetes Service (Amazon EKS) with an Amazon EC2 deployment to deploy a containerized application. Use an Amazon Aurora cluster to store the product data. Use AWS Step Functions to manage workflows. Store the product images in Amazon S3 Glacier Deep Archive.
- D. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate to deploy a containerized application. Use Amazon RDS with a Multi-AZ deployment to store the product data. Store the product images in an Amazon S3 bucket.

Answer: D

**QUESTION 1064** A company is migrating an application from an on-premises environment to AWS. The application will store sensitive data in Amazon S3. The company must encrypt the data before storing the data in Amazon S3. Which solution will meet these requirements?

- A. Encrypt the data by using client-side encryption with customer managed keys.
- B. Encrypt the data by using server-side encryption with AWS KMS keys (SSE-KMS).
- C. Encrypt the data by using server-side encryption with customer-provided keys (SSE-C).
- D. Encrypt the data by using client-side encryption with Amazon S3 managed keys.

Answer: A

Explanation: If client-side encryption is used, the keys must be managed by the customer.

**QUESTION 1065** A company wants to create an Amazon EMR cluster that multiple teams will use. The company wants to ensure that each team's big data workloads can access only the AWS services that each team needs to interact with. The company does not want the workloads to have access to Instance Metadata Service Version 2 (IMDSv2) on the cluster's underlying EC2 instances. Which solution will meet these requirements?

- A. Configure interface VPC endpoints for each AWS service that the teams need. Use the required interface VPC endpoints to submit the big data workloads.
- B. Create EMR runtime roles. Configure the cluster to use the runtime roles. Use the runtime roles to submit the big data workloads.
- C. Create an EC2 IAM instance profile that has the required permissions for each team. Use the instance profile to submit the big data workloads.
- D. Create an EMR security configuration that has the EnableApplicationScopedIAMRole option set to false. Use the security configuration to submit the big data workloads.

Answer: B

Explanation: EMR runtime roles allow fine-grained permissions per job, letting each team access only the services they are authorized to use. This isolates IAM permissions per workload and avoids exposing instance-level credentials through IMDSv2. Runtime roles improve security posture in multi-tenant EMR environments.

**QUESTION 1066** A solutions architect is designing an application that helps users fill out and submit registration forms. The solutions architect plans to use a two-tier architecture that includes a web application server tier and a worker tier. The application needs to process submitted forms quickly. The application needs to process each form exactly once. The solution must ensure that no data is lost. Which solution will meet these requirements?

- A. Use an Amazon Simple Queue Service (Amazon SQS) FIFO queue between the web application server tier and the worker tier to store and forward form data.
- B. Use an Amazon API Gateway HTTP API between the web application server tier and the worker tier to store and forward form data.
- C. Use an Amazon Simple Queue Service (Amazon SQS) standard queue between the web application server tier and the worker tier to store and forward form data.
- D. Use an AWS Step Functions

workflow. Create a synchronous workflow between the web application server tier and the worker tier that stores and forwards form data.

Answer: AQUESTION 1067A finance company uses an on-premises search application to collect streaming data from various producers. The application provides real-time updates to search and visualization features. The company is planning to migrate to AWS and wants to use an AWS native solution. Which solution will meet these requirements?

A. Use Amazon EC2 instances to ingest and process the data streams to Amazon S3 buckets for storage. Use Amazon Athena to search the data. Use Amazon Managed Grafana to create visualizations.

B. Use Amazon EMR to ingest and process the data streams to Amazon Redshift for storage. Use Amazon Redshift Spectrum to search the data. Use Amazon QuickSight to create visualizations.

C. Use Amazon Elastic Kubernetes Service (Amazon EKS) to ingest and process the data streams to Amazon DynamoDB for storage. Use Amazon CloudWatch to create graphical dashboards to search and visualize the data.

D. Use Amazon Kinesis Data Streams to ingest and process the data streams to Amazon OpenSearch Service. Use OpenSearch Service to search the data. Use Amazon QuickSight to create visualizations.

Answer: DQUESTION 1068A company currently runs an on-premises application that uses ASP.NET on Linux machines. The application is resource-intensive and serves customers directly. The company wants to modernize the application to .NET. The company wants to run the application on containers and to scale based on Amazon CloudWatch metrics. The company also wants to reduce the time spent on operational maintenance activities. Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS App2Container to containerize the application. Use an AWS CloudFormation template to deploy the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate.

B. Use AWS App2Container to containerize the application. Use an AWS CloudFormation template to deploy the application to Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 instances.

C. Use AWS App Runner to containerize the application. Use App Runner to deploy the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate.

D. Use AWS App Runner to containerize the application. Use App Runner to deploy the application to Amazon Elastic Kubernetes Service (Amazon EKS) on Amazon EC2 instances.

Answer: AQUESTION 1069A company is designing a new internal web application in the AWS Cloud. The new application must securely retrieve and store multiple employee usernames and passwords from an AWS managed service. Which solution will meet these requirements with the LEAST operational overhead?

A. Store the employee credentials in AWS Systems Manager Parameter Store. Use AWS CloudFormation and the BatchGetSecretValue API to retrieve usernames and passwords from Parameter Store.

B. Store the employee credentials in AWS Secrets Manager. Use AWS CloudFormation and AWS Batch with the BatchGetSecretValue API to retrieve the usernames and passwords from Secrets Manager.

C. Store the employee credentials in AWS Systems Manager Parameter Store. Use AWS CloudFormation and AWS Batch with the BatchGetSecretValue API to retrieve the usernames and passwords from Parameter Store.

D. Store the employee credentials in AWS Secrets Manager. Use AWS CloudFormation and the BatchGetSecretValue API to retrieve the usernames and passwords from Secrets Manager.

Answer: DQUESTION 1070A company that is in the ap-northeast-1 Region has a fleet of thousands of AWS Outposts servers. The company has deployed the servers at remote locations around the world. All the servers regularly download new software versions that consist of 100 files. There is significant latency before all servers run the new software versions. The company must reduce the deployment latency for new software versions. Which solution will meet this requirement with the LEAST operational overhead?

A. Create an Amazon S3 bucket in ap-northeast-1. Set up an Amazon CloudFront distribution in ap-northeast-1 that includes a CachingDisabled cache policy. Configure the S3 bucket as the origin. Download the software by using signed URLs.

B. Create an Amazon S3 bucket in ap-northeast-1. Create a second S3 bucket in the us-east-1 Region. Configure replication between the buckets. Set up an Amazon CloudFront distribution that uses ap-northeast-1 as the primary origin and us-east-1 as the secondary origin. Download the software by using signed URLs.

C. Create an Amazon S3 bucket in ap-northeast-1. Configure Amazon S3 Transfer Acceleration. Download the software by using the S3 Transfer Acceleration endpoint.

D. Create an Amazon S3 bucket in ap-northeast-1. Set up an Amazon CloudFront distribution. Configure the S3 bucket as the origin. Download the software by using signed URLs.

Answer: DQUESTION 1071A company currently runs an on-premises stock trading application by using Microsoft Windows Server. The company wants to migrate the application to the AWS Cloud. The company needs to design a highly available solution that provides low-latency access to block storage across multiple Availability Zones. Which solution will meet these requirements with the LEAST implementation effort?

A. Configure a Windows Server cluster that spans two Availability Zones on Amazon EC2 instances. Install the application on both cluster nodes. Use Amazon FSx for Windows File Server as shared storage between the two cluster nodes.

B. Configure a Windows Server cluster that spans two Availability Zones on Amazon EC2 instances. Install the application on both cluster nodes. Use Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp3) volumes as storage attached to the EC2 instances. Set up application-level replication to sync data from one EBS volume in one Availability Zone to another EBS volume in the second Availability Zone.

C. Deploy the application on Amazon EC2 instances in two Availability Zones. Configure one EC2 instance as active and the second EC2 instance in standby mode. Use an Amazon FSx for NetApp ONTAP Multi-AZ file

system to access the data by using Internet Small Computer Systems Interface (iSCSI) protocol.D. Deploy the application on Amazon EC2 instances in two Availability Zones. Configure one EC2 instance as active and the second EC2 instance in standby mode. Use Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io2) volumes as storage attached to the EC2 instances. Set up Amazon EBS level replication to sync data from one io2 volume in one Availability Zone to another io2 volume in the second Availability Zone.

Answer: AQUESTION 1072A company is designing a web application with an internet-facing Application Load Balancer (ALB). The company needs the ALB to receive HTTPS web traffic from the public internet. The ALB must send only HTTPS traffic to the web application servers hosted on the Amazon EC2 instances on port 443. The ALB must perform a health check of the web application servers over HTTPS on port 8443. Which combination of configurations of the security group that is associated with the ALB will meet these requirements? (Choose three.)

- A. Allow HTTPS inbound traffic from 0.0.0.0/0 for port 443.
- B. Allow all outbound traffic to 0.0.0.0/0 for port 443.
- C. Allow HTTPS outbound traffic to the web application instances for port 443.
- D. Allow HTTPS inbound traffic from the web application instances for port 443.
- E. Allow HTTPS outbound traffic to the web application instances for the health check on port 8443.
- F. Allow HTTPS inbound traffic from the web application instances for the health check on port 8443.

Answer: ACEQUESTION 1073A company hosts an application on AWS. The application gives users the ability to upload photos and store the photos in an Amazon S3 bucket. The company wants to use Amazon CloudFront and a custom domain name to upload the photo files to the S3 bucket in the eu-west-1 Region. Which solution will meet these requirements? (Choose two.)

- A. Use AWS Certificate Manager (ACM) to create a public certificate in the us-east-1 Region. Use the certificate in CloudFront.
- B. Use AWS Certificate Manager (ACM) to create a public certificate in eu-west-1. Use the certificate in CloudFront.
- C. Configure Amazon S3 to allow uploads from CloudFront. Configure S3 Transfer Acceleration.
- D. Configure Amazon S3 to allow uploads from CloudFront origin access control (OAC).
- E. Configure Amazon S3 to allow uploads from CloudFront. Configure an Amazon S3 website endpoint.

Answer: ADQUESTION 1074A weather forecasting company collects temperature readings from various sensors on a continuous basis. An existing data ingestion process collects the readings and aggregates the readings into larger Apache Parquet files. Then the process encrypts the files by using client-side encryption with KMS managed keys (CSE-KMS). Finally, the process writes the files to an Amazon S3 bucket with separate prefixes for each calendar day. The company wants to run occasional SQL queries on the data to take sample moving averages for a specific calendar day. Which solution will meet these requirements MOST cost-effectively?

- A. Configure Amazon Athena to read the encrypted files. Run SQL queries on the data directly in Amazon S3.
- B. Use Amazon S3 Select to run SQL queries on the data directly in Amazon S3.
- C. Configure Amazon Redshift to read the encrypted files. Use Redshift Spectrum and Redshift query editor v2 to run SQL queries on the data directly in Amazon S3.
- D. Configure Amazon EMR Serverless to read the encrypted files. Use Apache SparkSQL to run SQL queries on the data directly in Amazon S3.

Answer: AQUESTION 1075A company is implementing a new application on AWS. The company will run the application on multiple Amazon EC2 instances across multiple Availability Zones within multiple AWS Regions. The application will be available through the internet. Users will access the application from around the world. The company wants to ensure that each user who accesses the application is sent to the EC2 instances that are closest to the user's location. Which solution will meet these requirements?

- A. Implement an Amazon Route 53 geolocation routing policy. Use an internet-facing Application Load Balancer to distribute the traffic across all Availability Zones within the same Region.
- B. Implement an Amazon Route 53 geoproximity routing policy. Use an internet-facing Network Load Balancer to distribute the traffic across all Availability Zones within the same Region.
- C. Implement an Amazon Route 53 multivalue answer routing policy. Use an internet-facing Application Load Balancer to distribute the traffic across all Availability Zones within the same Region.
- D. Implement an Amazon Route 53 weighted routing policy. Use an internet-facing Network Load Balancer to distribute the traffic across all Availability Zones within the same Region.

Answer: BQUESTION 1076A financial services company plans to launch a new application on AWS to handle sensitive financial transactions. The company will deploy the application on Amazon EC2 instances. The company will use Amazon RDS for MySQL as the database. The company's security policies mandate that data must be encrypted at rest and in transit. Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure encryption at rest for Amazon RDS for MySQL by using AWS KMS managed keys. Configure AWS Certificate Manager (ACM) SSL/TLS certificates for encryption in transit.
- B. Configure encryption at rest for Amazon RDS for MySQL by using AWS KMS managed keys. Configure IPsec tunnels for encryption in transit.
- C. Implement third-party application-level data encryption before storing data in Amazon RDS for MySQL. Configure AWS Certificate Manager (ACM) SSL/TLS certificates for encryption in transit.
- D. Configure encryption at rest for Amazon RDS for MySQL by using AWS KMS managed keys. Configure a VPN connection to enable private connectivity to encrypt data in transit.

Answer: AQUESTION 1077A company is migrating its on-premises Oracle database to an Amazon RDS for Oracle database. The company needs to retain data for 90 days to meet regulatory requirements. The company must also be able to restore the database to a specific point in time

for up to 14 days. Which solution will meet these requirements with the LEAST operational overhead? A. Create Amazon RDS automated backups. Set the retention period to 90 days. B. Create an Amazon RDS manual snapshot every day. Delete manual snapshots that are older than 90 days. C. Use the Amazon Aurora Clone feature for Oracle to create a point-in-time restore. Delete clones that are older than 90 days. D. Create a backup plan that has a retention period of 90 days by using AWS Backup for Amazon RDS.

**Answer:** D

QUESTION 1078

A company is developing a new application that uses a relational database to store user data and application configurations. The company expects the application to have steady user growth. The company expects the database usage to be variable and read-heavy, with occasional writes. The company wants to cost-optimize the database solution. The company wants to use an AWS managed database solution that will provide the necessary performance. Which solution will meet these requirements MOST cost-effectively?

A. Deploy the database on Amazon RDS. Use Provisioned IOPS SSD storage to ensure consistent performance for read and write operations.

B. Deploy the database on Amazon Aurora Serverless to automatically scale the database capacity based on actual usage to accommodate the workload.

C. Deploy the database on Amazon DynamoDB. Use on-demand capacity mode to automatically scale throughput to accommodate the workload.

D. Deploy the database on Amazon RDS. Use magnetic storage and use read replicas to accommodate the workload.

**Answer:** B

QUESTION 1079

A company hosts its application on several Amazon EC2 instances inside a VPC. The company creates a dedicated Amazon S3 bucket for each customer to store their relevant information in Amazon S3. The company wants to ensure that the application running on EC2 instances can securely access only the S3 buckets that belong to the company's AWS account. Which solution will meet these requirements with the LEAST operational overhead?

A. Create a gateway endpoint for Amazon S3 that is attached to the VPC. Update the IAM instance profile policy to provide access to only the specific buckets that the application needs.

B. Create a NAT gateway in a public subnet with a security group that allows access to only Amazon S3. Update the route tables to use the NAT Gateway.

C. Create a gateway endpoint for Amazon S3 that is attached to the VPC. Update the IAM instance profile policy with a Deny action and the following condition key:

D. Create a NAT Gateway in a public subnet. Update route tables to use the NAT Gateway. Assign bucket policies for all buckets with a Deny action and the following condition key:

**Answer:** C

QUESTION 1080

A company is building a cloud-based application on AWS that will handle sensitive customer data. The application uses Amazon RDS for the database, Amazon S3 for object storage, and S3 Event Notifications that invoke AWS Lambda for serverless processing. The company uses AWS IAM Identity Center to manage user credentials. The development, testing, and operations teams need secure access to Amazon RDS and Amazon S3 while ensuring the confidentiality of sensitive customer data. The solution must comply with the principle of least privilege. Which solution meets these requirements with the LEAST operational overhead?

A. Use IAM roles with least privilege to grant all the teams access. Assign IAM roles to each team with customized IAM policies defining specific permission for Amazon RDS and S3 object access based on team responsibilities.

B. Enable IAM Identity Center with an Identity Center directory. Create and configure permission sets with granular access to Amazon RDS and Amazon S3. Assign all the teams to groups that have specific access with the permission sets.

C. Create individual IAM users for each member in all the teams with role-based permissions. Assign the IAM roles with predefined policies for RDS and S3 access to each user based on user needs. Implement IAM Access Analyzer for periodic credential evaluation.

D. Use AWS Organizations to create separate accounts for each team. Implement cross-account IAM roles with least privilege. Grant specific permission for RDS and S3 access based on team roles and responsibilities.

**Answer:** B

**Explanation:** IAM Identity Center: This service simplifies user management by centralizing credentials and access control. Permission Sets: You can create granular permission sets that align with the principle of least privilege, ensuring that each team has only the access they need. Group Assignments: By assigning teams to groups with specific permission sets, you streamline access management and reduce the complexity of individual user permissions. This approach minimizes operational overhead while maintaining secure and compliant access to sensitive customer data.

QUESTION 1081

A company has an Amazon S3 bucket that contains sensitive data files. The company has an application that runs on virtual machines in an on-premises data center. The company currently uses AWS IAM Identity Center. The application requires temporary access to files in the S3 bucket. The company wants to grant the application secure access to the files in the S3 bucket. Which solution will meet these requirements?

A. Create an S3 bucket policy that permits access to the bucket from the public IP address range of the company's on-premises data center.

B. Use IAM Roles Anywhere to obtain security credentials in IAM Identity Center that grant access to the S3 bucket. Configure the virtual machines to assume the role by using the AWS CLI.

C. Install the AWS CLI on the virtual machine. Configure the AWS CLI with access keys from an IAM user that has access to the bucket.

D. Create an IAM user and policy that grants access to the bucket. Store the access key and secret key for the IAM user in AWS Secrets Manager. Configure the application to retrieve the access key and secret key at startup.

**Answer:** B

QUESTION 1082

A company hosts its core network services, including directory services and DNS, in its on-premises data center. The data center is connected to the AWS Cloud using AWS Direct Connect (DX). Additional AWS accounts are planned that will require quick, cost-effective, and consistent

access to these network services. What should a solutions architect implement to meet these requirements with the LEAST amount of operational overhead?A. Create a DX connection in each new account. Route the network traffic to the on-premises servers.B. Configure VPC endpoints in the DX VPC for all required services. Route the network traffic to the on-premises servers.C. Create a VPN connection between each new account and the DX VPC. Route the network traffic to the on-premises servers.D. Configure AWS Transit Gateway between the accounts. Assign DX to the transit gateway and route network traffic to the on-premises servers.

Answer: D

QUESTION 1083

A company hosts its main public web application in one AWS Region across multiple Availability Zones. The application uses an Amazon EC2 Auto Scaling group and an Application Load Balancer (ALB). A web development team needs a cost-optimized compute solution to improve the company's ability to serve dynamic content globally to millions of customers. Which solution will meet these requirements?A. Create an Amazon CloudFront distribution. Configure the existing ALB as the origin.B. Use Amazon Route 53 to serve traffic to the ALB and EC2 instances based on the geographic location of each customer.C. Create an Amazon S3 bucket with public read access enabled. Migrate the web application to the S3 bucket. Configure the S3 bucket for website hosting.D. Use AWS Direct Connect to directly serve content from the web application to the location of each customer.

Answer: A

QUESTION 1084

A company stores user data in AWS. The data is used continuously with peak usage during business hours. Access patterns vary, with some data not being used for months at a time. A solutions architect must choose a cost-effective solution that maintains the highest level of durability while maintaining high availability. Which storage solution meets these requirements?A. Amazon S3 StandardB. Amazon S3 Intelligent-TieringC. Amazon S3 Glacier Deep ArchiveD. Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: B

QUESTION 1085

A company is testing an application that runs on an Amazon EC2 Linux instance. A single 500 GB Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2) volume is attached to the EC2 instance. The company will deploy the application on multiple EC2 instances in an Auto Scaling group. All instances require access to the data that is stored in the EBS volume. The company needs a highly available and resilient solution that does not introduce significant changes to the application's code. Which solution will meet these requirements?A. Provision an EC2 instance that uses NFS server software. Attach a single 500 GB gp2 EBS volume to the instance.B. Provision an Amazon FSx for Windows File Server file system. Configure the file system as an SMB file store within a single Availability Zone.C. Provision an EC2 instance with two 250 GB Provisioned IOPS SSD EBS volumes.D. Provision an Amazon Elastic File System (Amazon EFS) file system. Configure the file system to use General Purpose performance mode.

Answer: D

Explanation: Amazon EFS is a fully managed, scalable file storage service that can be accessed concurrently by thousands of EC2 instances. It supports General Purpose performance mode for latency-sensitive use cases and provides high availability and durability across multiple Availability Zones with minimal changes to application code.

QUESTION 1086

A company recently launched a new application for its customers. The application runs on multiple Amazon EC2 instances across two Availability Zones. End users use TCP to communicate with the application. The application must be highly available and must automatically scale as the number of users increases. Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

A. Add a Network Load Balancer in front of the EC2 instances.B. Configure an Auto Scaling group for the EC2 instances.C. Add an Application Load Balancer in front of the EC2 instances.D. Manually add more EC2 instances for the application.E. Add a Gateway Load Balancer in front of the EC2 instances.

Answer: AB

Explanation: For an application requiring TCP communication and high availability, Network Load Balancer (NLB) is the best choice for load balancing TCP traffic because it is designed for handling high-throughput, low-latency connections. Auto Scaling group ensures that the application can automatically scale based on demand, adding or removing EC2 instances as needed, which is crucial for handling user growth.

QUESTION 1087

A company is designing the architecture for a new mobile app that uses the AWS Cloud. The company uses organizational units (OUs) in AWS Organizations to manage its accounts. The company wants to tag Amazon EC2 instances with data sensitivity by using values of sensitive and nonsensitive. IAM identities must not be able to delete a tag or create instances without a tag. Which combination of steps will meet these requirements? (Choose two.)

A. In Organizations, create a new tag policy that specifies the data sensitivity tag key and the required values. Enforce the tag values for the EC2 instances. Attach the tag policy to the appropriate OU.B. In Organizations, create a new service control policy (SCP) that specifies the data sensitivity tag key and the required tag values. Enforce the tag values for the EC2 instances. Attach the SCP to the appropriate OU.C. Create a tag policy to deny running instances when a tag key is not specified. Create another tag policy that prevents identities from deleting tags. Attach the tag policies to the appropriate OU.D. Create a service control policy (SCP) to deny creating instances when a tag key is not specified. Create another SCP that prevents identities from deleting tags. Attach the SCPs to the appropriate OU.E. Create an AWS Config rule to check if EC2 instances use the data sensitivity tag and the specified values. Configure an AWS Lambda function to delete the resource if a noncompliant resource is found.

Answer: AD

QUESTION 1088

A company runs database workloads on AWS that are the backend for the company's customer portals. The company runs a Multi-AZ database

cluster on Amazon RDS for PostgreSQL. The company needs to implement a 30-day backup retention policy. The company currently has both automated RDS backups and manual RDS backups. The company wants to maintain both types of existing RDS backups that are less than 30 days old. Which solution will meet these requirements MOST cost-effectively?A. Configure the RDS backup retention policy to 30 days for automated backups by using AWS Backup. Manually delete manual backups that are older than 30 days.B. Disable RDS automated backups. Delete automated backups and manual backups that are older than 30 days. Configure the RDS backup retention policy to 30 days for automated backups.C. Configure the RDS backup retention policy to 30 days for automated backups. Manually delete manual backups that are older than 30 days.D. Disable RDS automated backups. Delete automated backups and manual backups that are older than 30 days automatically by using AWS CloudFormation. Configure the RDS backup retention policy to 30 days for automated backups.

Answer: C

QUESTION 1089A company is planning to migrate a legacy application to AWS. The application currently uses NFS to communicate to an on-premises storage solution to store application data. The application cannot be modified to use any other communication protocols other than NFS for this purpose. Which storage solution should a solutions architect recommend for use after the migration?A. AWS DataSyncB. Amazon Elastic Block Store (Amazon EBS)C. Amazon Elastic File System (Amazon EFS)D. Amazon EMR File System (Amazon EMRFS)

Answer: C

QUESTION 1090A company uses GPS trackers to document the migration patterns of thousands of sea turtles. The trackers check every 5 minutes to see if a turtle has moved more than 100 yards (91.4 meters). If a turtle has moved, its tracker sends the new coordinates to a web application running on three Amazon EC2 instances that are in multiple Availability Zones in one AWS Region. Recently, the web application was overwhelmed while processing an unexpected volume of tracker data. Data was lost with no way to replay the events. A solutions architect must prevent this problem from happening again and needs a solution with the least operational overhead. What should the solutions architect do to meet these requirements?A. Create an Amazon S3 bucket to store the data. Configure the application to scan for new data in the bucket for processing.B. Create an Amazon API Gateway endpoint to handle transmitted location coordinates. Use an AWS Lambda function to process each item concurrently.C. Create an Amazon Simple Queue Service (Amazon SQS) queue to store the incoming data. Configure the application to poll for new messages for processing.D. Create an Amazon DynamoDB table to store transmitted location coordinates. Configure the application to query the table for new data for processing. Use TTL to remove data that has been processed.

Answer: C

QUESTION 1091A company's software development team needs an Amazon RDS Multi-AZ cluster. The RDS cluster will serve as a backend for a desktop client that is deployed on premises. The desktop client requires direct connectivity to the RDS cluster. The company must give the development team the ability to connect to the cluster by using the client when the team is in the office. Which solution provides the required connectivity MOST securely?A. Create a VPC and two public subnets. Create the RDS cluster in the public subnets. Use AWS Site-to-Site VPN with a customer gateway in the company's office.B. Create a VPC and two private subnets. Create the RDS cluster in the private subnets. Use AWS Site-to-Site VPN with a customer gateway in the company's office.C. Create a VPC and two private subnets. Create the RDS cluster in the private subnets. Use RDS security groups to allow the company's office IP ranges to access the cluster.D. Create a VPC and two public subnets. Create the RDS cluster in the public subnets. Create a cluster user for each developer. Use RDS security groups to allow the users to access the cluster.

Answer: C

QUESTION 1092A solutions architect is creating an application that will handle batch processing of large amounts of data. The input data will be held in Amazon S3 and the output data will be stored in a different S3 bucket. For processing, the application will transfer the data over the network between multiple Amazon EC2 instances. What should the solutions architect do to reduce the overall data transfer costs?A. Place all the EC2 instances in an Auto Scaling group.B. Place all the EC2 instances in the same AWS Region.C. Place all the EC2 instances in the same Availability Zone.D. Place all the EC2 instances in private subnets in multiple Availability Zones.

Answer: D

QUESTION 1093A company hosts a multi-tier web application that uses an Amazon Aurora MySQL DB cluster for storage. The application tier is hosted on Amazon EC2 instances. The company's IT security guidelines mandate that the database credentials be encrypted and rotated every 14 days. What should a solutions architect do to meet this requirement with the LEAST operational effort?A. Create a new AWS Key Management Service (AWS KMS) encryption key. Use AWS Secrets Manager to create a new secret that uses the KMS key with the appropriate credentials. Associate the secret with the Aurora DB cluster. Configure a custom rotation period of 14 days.B. Create two parameters in AWS Systems Manager Parameter Store: one for the user name as a string parameter and one that uses the SecureString type for the password. Select AWS Key Management Service (AWS KMS) encryption for the password parameter, and load these parameters in the application tier. Implement an AWS Lambda function that rotates the password every 14 days.C. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system in all EC2 instances of the application tier. Restrict the access to the file on the file system so that the application can read the file and that only super users can modify the file. Implement an AWS Lambda function that rotates the key in Aurora every 14 days.

14 days and writes new credentials into the file.D. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon S3 bucket that the application uses to load the credentials. Download the file to the application regularly to ensure that the correct credentials are used. Implement an AWS Lambda function that rotates the Aurora credentials every 14 days and uploads these credentials to the file in the S3 bucket.Answer: AQUESTION 1094A streaming media company is rebuilding its infrastructure to accommodate increasing demand for video content that users consume daily.The company needs to process terabyte-sized videos to block some content in the videos. Video processing can take up to 20 minutes.The company needs a solution that will scale with demand and remain cost-effective.Which solution will meet these requirements?A. Use AWS Lambda functions to process videos. Store video metadata in Amazon DynamoDB. Store video content in Amazon S3 Intelligent-Tiering.B. Use Amazon Elastic Container Service (Amazon ECS) and AWS Fargate to implement microservices to process videos. Store video metadata in Amazon Aurora. Store video content in Amazon S3 Intelligent-Tiering.C. Use Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB) to process videos. Store video content in Amazon S3 Standard. Use Amazon Simple Queue Service (Amazon SQS) for queuing and to decouple processing tasks.D. Deploy a containerized video processing application on Amazon Elastic Kubernetes Service (Amazon EKS) on Amazon EC2. Store video metadata in Amazon RDS in a single Availability Zone. Store video content in Amazon S3 Glacier Deep Archive.Answer: BQUESTION 1095A company runs an on-premises application on a Kubernetes cluster. The company recently added millions of new customers. The company's existing on-premises infrastructure is unable to handle the large number of new customers. The company needs to migrate the on-premises application to the AWS Cloud.The company will migrate to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The company does not want to manage the underlying compute infrastructure for the new architecture on AWS.Which solution will meet these requirements with the LEAST operational overhead?A. Use a self-managed node to supply compute capacity. Deploy the application to the new EKS cluster.B. Use managed node groups to supply compute capacity. Deploy the application to the new EKS cluster.C. Use AWS Fargate to supply compute capacity. Create a Fargate profile. Use the Fargate profile to deploy the application.D. Use managed node groups with Karpenter to supply compute capacity. Deploy the application to the new EKS cluster.Answer: CQUESTION 1096A company is launching a new application that requires a structured database to store user profiles, application settings, and transactional data. The database must be scalable with application traffic and must offer backups.Which solution will meet these requirements MOST cost-effectively?A. Deploy a self-managed database on Amazon EC2 instances by using open source software. Use Spot Instances for cost optimization. Configure automated backups to Amazon S3.B. Use Amazon RDS. Use on-demand capacity mode for the database with General Purpose SSD storage. Configure automatic backups with a retention period of 7 days.C. Use Amazon Aurora Serverless for the database. Use serverless capacity scaling. Configure automated backups to Amazon S3.D. Deploy a self-managed NoSQL database on Amazon EC2 instances. Use Reserved Instances for cost optimization. Configure automated backups directly to Amazon S3 Glacier Flexible Retrieval.Answer: C Explanation:Amazon Aurora Serverless v2 is ideal for applications with unpredictable or intermittent workloads. It automatically scales capacity up or down based on demand, significantly reducing costs. It supports automated backups to Amazon S3, making it suitable and cost-effective for new applications with variable traffic.QUESTION 1097A company runs its legacy web application on AWS. The web application server runs on an Amazon EC2 instance in the public subnet of a VPC. The web application server collects images from customers and stores the image files in a locally attached Amazon Elastic Block Store (Amazon EBS) volume. The image files are uploaded every night to an Amazon S3 bucket for backup.A solutions architect discovers that the image files are being uploaded to Amazon S3 through the public endpoint. The solutions architect needs to ensure that traffic to Amazon S3 does not use the public endpoint.Which solution will meet these requirements?A. Create a gateway VPC endpoint for the S3 bucket that has the necessary permissions for the VPC. Configure the subnet route table to use the gateway VPC endpoint.B. Move the S3 bucket inside the VPC. Configure the subnet route table to access the S3 bucket through private IP addresses.C. Create an Amazon S3 access point for the Amazon EC2 instance inside the VPCConfigure the web application to upload by using the Amazon S3 access point.D. Configure an AWS Direct Connect connection between the VPC that has the Amazon EC2 instance and Amazon S3 to provide a dedicated network path.Answer: AQUESTION 1098A company is creating a prototype of an ecommerce website on AWS. The website consists of an Application Load Balancer, an Auto Scaling group of Amazon EC2 instances for web servers, and an Amazon RDS for MySQL DB instance that runs with the Single-AZ configuration.The website is slow to respond during searches of the product catalog. The product catalog is a group of tables in the MySQL database that the company does not update frequently. A solutions architect has determined that the CPU utilization on the DB instance is high when product catalog searches occur.What should the solutions architect recommend to improve the performance of the website during searches of the product catalog?A. Migrate the product catalog to an Amazon Redshift database. Use the COPY command to load the product catalog tables.B. Implement an Amazon ElastiCache for Redis cluster to cache the product catalog. Use lazy loading to populate the cache.

C. Add an additional scaling policy to the Auto Scaling group to launch additional EC2 instances when database response is slow.

D. Turn on the Multi-AZ configuration for the DB instance. Configure the EC2 instances to throttle the product catalog queries that are sent to the database.

Answer: B

QUESTION 1099A company currently stores 5 TB of data in on-premises block storage systems. The company's current storage solution provides limited space for additional data. The company runs applications on premises that must be able to retrieve frequently accessed data with low latency. The company requires a cloud-based storage solution. Which solution will meet these requirements with the MOST operational efficiency?

A. Use Amazon S3 File Gateway. Integrate S3 File Gateway with the on-premises applications to store and directly retrieve files by using the SMB file system.

B. Use an AWS Storage Gateway Volume Gateway with cached volumes as iSCSI targets.

C. Use an AWS Storage Gateway Volume Gateway with stored volumes as iSCSI targets.

D. Use an AWS Storage Gateway Tape Gateway. Integrate Tape Gateway with the on-premises applications to store virtual tapes in Amazon S3.

Answer: B

QUESTION 1100A company operates a food delivery service. Because of recent growth, the company's order processing system is experiencing scaling problems during peak traffic hours. The current architecture includes Amazon EC2 instances in an Auto Scaling group that collect orders from an application. A second group of EC2 instances in an Auto Scaling group fulfills the 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 adequately during peak traffic hours. Which solution will meet these requirements?

A. Use Amazon CloudWatch to monitor the CPUUtilization metric for each instance in both Auto Scaling groups. Configure each Auto Scaling group's minimum capacity to meet its peak workload value.

B. Use Amazon CloudWatch to monitor the CPUUtilization metric for each instance in both Auto Scaling groups. Configure a CloudWatch alarm to invoke an Amazon Simple Notification Service (Amazon SNS) topic to create additional Auto Scaling groups on demand.

C. Provision two Amazon Simple Queue Service (Amazon SQS) queues. Use one SQS queue for order collection. Use the second SQS queue for order fulfillment. Configure the EC2 instances to poll their respective queues. Scale the Auto Scaling groups based on notifications that the queues send.

D. Provision two Amazon Simple Queue Service (Amazon SQS) queues. Use one SQS queue for order collection. Use the second SQS queue for order fulfillment. Configure the EC2 instances to poll their respective queues. Scale the Auto Scaling groups based on the number of messages in each queue.

Answer: D

QUESTION 1101An online gaming company is transitioning user data storage to Amazon DynamoDB to support the company's growing user base. The current architecture includes DynamoDB tables that contain user profiles, achievements, and in-game transactions. The company needs to design a robust, continuously available, and resilient DynamoDB architecture to maintain a seamless gaming experience for users. Which solution will meet these requirements MOST cost-effectively?

A. Create DynamoDB tables in a single AWS Region. Use on-demand capacity mode. Use global tables to replicate data across multiple Regions.

B. Use DynamoDB Accelerator (DAX) to cache frequently accessed data. Deploy tables in a single AWS Region and enable auto scaling. Configure Cross-Region Replication manually to additional Regions.

C. Create DynamoDB tables in multiple AWS Regions. Use on-demand capacity mode. Use DynamoDB Streams for Cross-Region Replication between Regions.

D. Use DynamoDB global tables for automatic multi-Region replication. Deploy tables in multiple AWS Regions. Use provisioned capacity mode. Enable auto scaling.

Answer: D

Explanation: DynamoDB Global Tables provide a fully managed, multi-region, and multi-master database solution that allows you to deploy DynamoDB tables in multiple AWS Regions. This ensures high availability and resiliency across different geographical locations, providing a seamless gaming experience for users. Using provisioned capacity mode with auto-scaling ensures cost-efficiency by scaling up or down based on actual demand.

QUESTION 1102A company runs its media rendering application on-premises. The company wants to reduce storage costs and has moved all data to Amazon S3. The on-premises rendering application needs low-latency access to storage. The company needs to design a storage solution for the application. The storage solution must maintain the desired application performance. Which storage solution will meet these requirements in the MOST cost-effective way?

A. Use Mountpoint for Amazon S3 to access the data in Amazon S3 for the on-premises application.

B. Configure an Amazon S3 File Gateway to provide storage for the on-premises application.

C. Copy the data from Amazon S3 to Amazon FSx for Windows File Server. Configure an Amazon FSx File Gateway to provide storage for the on-premises application.

D. Configure an on-premises file server. Use the Amazon S3 API to connect to S3 storage. Configure the application to access the storage from the on-premises file server.

Answer: B

QUESTION 1103A company hosts its enterprise resource planning (ERP) system in the us-east-1 Region. The system runs on Amazon EC2 instances. Customers use a public API that is hosted on the EC2 instances to exchange information with the ERP system. International customers report slow API response times from their data centers. Which solution will improve response times for the international customers MOST cost-effectively?

A. Create an AWS Direct Connect connection that has a public virtual interface (VIF) to provide connectivity from each customer's data center to us-east-1. Route customer API requests by using a Direct Connect gateway to the ERP system API.

B. Set up an Amazon CloudFront distribution in

front of the API. Configure the CachingOptimized managed cache policy to provide improved cache efficiency.C. Set up AWS Global Accelerator. Configure listeners for the necessary ports. Configure endpoint groups for the appropriate Regions to distribute traffic. Create an endpoint in the group for the API.D. Use AWS Site-to-Site VPN to establish dedicated VPN tunnels between Regions and customer networks. Route traffic to the API over the VPN connections.

Answer: B

QUESTION 1104A company tracks customer satisfaction by using surveys that the company hosts on its website. The surveys sometimes reach thousands of customers every hour. Survey results are currently sent in email messages to the company so company employees can manually review results and assess customer sentiment. The company wants to automate the customer survey process. Survey results must be available for the previous 12 months. Which solution will meet these requirements in the MOST scalable way?

A. Send the survey results data to an Amazon API Gateway endpoint that is connected to an Amazon Simple Queue Service (Amazon SQS) queue. Create an AWS Lambda function to poll the SQS queue, call Amazon Comprehend for sentiment analysis, and save the results to an Amazon DynamoDB table. Set the TTL for all records to 365 days in the future.

B. Send the survey results data to an API that is running on an Amazon EC2 instance. Configure the API to store the survey results as a new record in an Amazon DynamoDB table, call Amazon Comprehend for sentiment analysis, and save the results in a second DynamoDB table. Set the TTL for all records to 365 days in the future.

C. Write the survey results data to an Amazon S3 bucket. Use S3 Event Notifications to invoke an AWS Lambda function to read the data and call Amazon Rekognition for sentiment analysis. Store the sentiment analysis results in a second S3 bucket. Use S3 lifecycle policies on each bucket to expire objects after 365 days.

D. Send the survey results data to an Amazon API Gateway endpoint that is connected to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the SQS queue to invoke an AWS Lambda function that calls Amazon Lex for sentiment analysis and saves the results to an Amazon DynamoDB table. Set the TTL for all records to 365 days in the future.

Answer: A

Explanation: This solution is the most scalable and efficient way to handle large volumes of survey data while automating sentiment analysis: API Gateway and SQS: The survey results are sent to API Gateway, which forwards the data to an SQS queue. SQS can handle large volumes of messages and ensures that messages are not lost. AWS Lambda: Lambda is triggered by polling the SQS queue, where it processes the survey data. Amazon Comprehend: Comprehend is used for sentiment analysis, providing insights into customer satisfaction. DynamoDB with TTL: Results are stored in DynamoDB with a Time to Live (TTL) attribute set to expire after 365 days, automatically removing old data and reducing storage costs.

Option B (EC2 API): Running an API on EC2 requires more maintenance and scalability management compared to API Gateway.

QUESTION 1105A company uses AWS Systems Manager for routine management and patching of Amazon EC2 instances. The EC2 instances are in an IP address type target group behind an Application Load Balancer (ALB). New security protocols require the company to remove EC2 instances from service during a patch. When the company attempts to follow the security protocol during the next patch, the company receives errors during the patching window. Which combination of solutions will resolve the errors? (Choose two.)

A. Change the target type of the target group from IP address type to instance type.

B. Continue to use the existing Systems Manager document without changes because it is already optimized to handle instances that are in an IP address type target group behind an ALB.

C. Implement the AWSEC2-PatchLoadBalancerInstance Systems Manager Automation document to manage the patching process.

D. Use Systems Manager Maintenance Windows to automatically remove the instances from service to patch the instances.

E. Configure Systems Manager State Manager to remove the instances from service and manage the patching schedule. Use ALB health checks to re-route traffic.

Answer: CD

QUESTION 1106A medical company wants to perform transformations on a large amount of clinical trial data that comes from several customers. The company must extract the data from a relational database that contains the customer data. Then the company will transform the data by using a series of complex rules. The company will load the data to Amazon S3 when the transformations are complete. All data must be encrypted where it is processed before the company stores the data in Amazon S3. All data must be encrypted by using customer-specific keys. Which solution will meet these requirements with the LEAST amount of operational effort?

A. Create one AWS Glue job for each customer. Attach a security configuration to each job that uses server-side encryption with Amazon S3 managed keys (SSE-S3) to encrypt the data.

B. Create one Amazon EMR cluster for each customer. Attach a security configuration to each cluster that uses client-side encryption with a custom client-side root key (CSE-Custom) to encrypt the data.

C. Create one AWS Glue job for each customer. Attach a security configuration to each job that uses client-side encryption with AWS KMS managed keys (CSE-KMS) to encrypt the data.

D. Create one Amazon EMR cluster for each customer. Attach a security configuration to each cluster that uses server-side encryption with AWS KMS keys (SSE-KMS) to encrypt the data.

Answer: C

QUESTION 1107A company hosts a website analytics application on a single Amazon EC2 On-Demand Instance. The analytics application is highly resilient and is designed to run in stateless mode. The company notices that the application is showing signs of performance degradation during busy times and is presenting 5xx errors. The company needs to make the application scale seamlessly. Which solution will meet these requirements MOST cost-effectively?

A. Create an Amazon Machine Image (AMI) of

the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use an Application Load Balancer to distribute the load across the two EC2 instances.B. Create an Amazon Machine Image (AMI) of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.C. Create an AWS Lambda function to stop the EC2 instance and change the instance type. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization is more than 75%.D. Create an Amazon Machine Image (AMI) of the web application. Apply the AMI to a launch template. Create an Auto Scaling group that includes the launch template. Configure the launch template to use a Spot Fleet. Attach an Application Load Balancer to the Auto Scaling group.Answer: DQUESTION 1108A company runs an environment where data is stored in an Amazon S3 bucket. The objects are accessed frequently throughout the day. The company has strict data encryption requirements for data that is stored in the S3 bucket. The company currently uses AWS Key Management Service (AWS KMS) for encryption. The company wants to optimize costs associated with encrypting S3 objects without making additional calls to AWS KMS. Which solution will meet these requirements?A. Use server-side encryption with Amazon S3 managed keys (SSE-S3).B. Use an S3 Bucket Key for server-side encryption with AWS KMS keys (SSE-KMS) on the new objects.C. Use client-side encryption with AWS KMS customer managed keys.D. Use server-side encryption with customer-provided keys (SSE-C) stored in AWS KMS.Answer: BQUESTION 1109A company runs multiple workloads on virtual machines (VMs) in an on-premises data center. The company is expanding rapidly. The on-premises data center is not able to scale fast enough to meet business needs. The company wants to migrate the workloads to AWS. The migration is time sensitive. The company wants to use a lift-and-shift strategy for non-critical workloads. Which combination of steps will meet these requirements? (Choose three.)A. Use the AWS Schema Conversion Tool (AWS SCT) to collect data about the VMs.B. Use AWS Application Migration Service. Install the AWS Replication Agent on the VMs.C. Complete the initial replication of the VMs. Launch test instances to perform acceptance tests on the VMs.D. Stop all operations on the VMs. Launch a cutover instance.E. Use AWS App2Container (A2C) to collect data about the VMs.F. Use AWS Database Migration Service (AWS DMS) to migrate the VMs.Answer: BCDQUESTION 1110A company hosts an application in a private subnet. The company has already integrated the application with Amazon Cognito. The company uses an Amazon Cognito user pool to authenticate users. The company needs to modify the application so the application can securely store user documents in an Amazon S3 bucket. Which combination of steps will securely integrate Amazon S3 with the application? (Choose two.)A. Create an Amazon Cognito identity pool to generate secure Amazon S3 access tokens for users when they successfully log in.B. Use the existing Amazon Cognito user pool to generate Amazon S3 access tokens for users when they successfully log in.C. Create an Amazon S3 VPC endpoint in the same VPC where the company hosts the application.D. Create a NAT gateway in the VPC where the company hosts the application. Assign a policy to the S3 bucket to deny any request that is not initiated from Amazon Cognito.E. Attach a policy to the S3 bucket that allows access only from the users' IP addresses.Answer: ACQUESTION 1111A company has a three-tier web application that processes orders from customers. The web tier consists of Amazon EC2 instances behind an Application Load Balancer. The processing tier consists of EC2 instances. The company decoupled the web tier and processing tier by using Amazon Simple Queue Service (Amazon SQS). The storage layer uses Amazon DynamoDB. At peak times, some users report order processing delays and hangs. The company has noticed that during these delays, the EC2 instances are running at 100% CPU usage, and the SQS queue fills up. The peak times are variable and unpredictable. The company needs to improve the performance of the application. Which solution will meet these requirements?A. Use scheduled scaling for Amazon EC2 Auto Scaling to scale out the processing tier instances for the duration of peak usage times. Use the CPU Utilization metric to determine when to scale.B. Use Amazon ElastiCache for Redis in front of the DynamoDB backend tier. Use target utilization as a metric to determine when to scale.C. Add an Amazon CloudFront distribution to cache the responses for the web tier. Use HTTP latency as a metric to determine when to scale.D. Use an Amazon EC2 Auto Scaling target tracking policy to scale out the processing tier instances. Use the ApproximateNumberOfMessages attribute to determine when to scale.Answer: DQUESTION 1112A company's production environment consists of Amazon EC2 On-Demand Instances that run constantly between Monday and Saturday. The instances must run for only 12 hours on Sunday and cannot tolerate interruptions. The company wants to cost-optimize the production environment. Which solution will meet these requirements MOST cost-effectively?A. Purchase Scheduled Reserved Instances for the EC2 instances that run for only 12 hours on Sunday. Purchase Standard Reserved Instances for the EC2 instances that run constantly between Monday and Saturday.B. Purchase Convertible Reserved Instances for the EC2 instances that run for only 12 hours on Sunday. Purchase Standard Reserved Instances for the EC2 instances that run constantly between Monday and Saturday.C. Use Spot Instances for the EC2 instances that run for only 12 hours on Sunday. Purchase Standard Reserved Instances for the EC2 instances that run constantly between Monday and Saturday.D. Use Spot Instances for the EC2 instances that run for only 12 hours on Sunday. Purchase Convertible Reserved Instances for the EC2 instances that run constantly between Monday and

Saturday.Answer: AQUESTION 1113A digital image processing company wants to migrate its on-premises monolithic application to the AWS Cloud. The company processes thousands of images and generates large files as part of the processing workflow. The company needs a solution to manage the growing number of image processing jobs. The solution must also reduce the manual tasks in the image processing workflow. The company does not want to manage the underlying infrastructure of the solution. Which solution will meet these requirements with the LEAST operational overhead?A. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 Spot Instances to process the images. Configure Amazon Simple Queue Service (Amazon SQS) to orchestrate the workflow. Store the processed files in Amazon Elastic File System (Amazon EFS).B. Use AWS Batch jobs to process the images. Use AWS Step Functions to orchestrate the workflow. Store the processed files in an Amazon S3 bucket.C. Use AWS Lambda functions and Amazon EC2 Spot Instances to process the images. Store the processed files in Amazon FSx.D. Deploy a group of Amazon EC2 instances to process the images. Use AWS Step Functions to orchestrate the workflow. Store the processed files in an Amazon Elastic Block Store (Amazon EBS) volume.Answer: BQUESTION 1114A company's image-hosting website gives users around the world the ability to up load, view, and download images from their mobile devices. The company currently hosts the static website in an Amazon S3 bucket. Because of the website's growing popularity, the website's performance has decreased. Users have reported latency issues when they upload and download images. The company must improve the performance of the website. Which solution will meet these requirements with the LEAST implementation effort?A. Configure an Amazon CloudFront distribution for the S3 bucket to improve the download performance. Enable S3 Transfer Acceleration to improve the upload performance.B. Configure Amazon EC2 instances of the right sizes in multiple AWS Regions. Migrate the application to the EC2 instances. Use an Application Load Balancer to distribute the website traffic equally among the EC2 instances. Configure AWS Global Accelerator to address global demand with low latency.C. Configure an Amazon CloudFront distribution that uses the S3 bucket as an origin to improve the download performance. Configure the application to use CloudFront to upload images to improve the upload performance. Create S3 buckets in multiple AWS Regions. Configure replication rules for the buckets to replicate users' data based on the users' location. Redirect downloads to the S3 bucket that is closest to each user's location.D. Configure AWS Global Accelerator for the S3 bucket to improve network performance. Create an endpoint for the application to use Global Accelerator instead of the S3 bucket.Answer: AQUESTION 1115A company runs an application in a private subnet behind an Application Load Balancer (ALB) in a VPC. The VPC has a NAT gateway and an internet gateway. The application calls the Amazon S3 API to store objects. According to the company's security policy, traffic from the application must not travel across the internet. Which solution will meet these requirements MOST cost-effectively?A. Configure an S3 interface endpoint. Create a security group that allows outbound traffic to Amazon S3.B. Configure an S3 gateway endpoint. Update the VPC route table to use the endpoint.C. Configure an S3 bucket policy to allow traffic from the Elastic IP address that is assigned to the NAT gateway.D. Create a second NAT gateway in the same subnet where the legacy application is deployed. Update the VPC route table to use the second NAT gateway.Answer: BQUESTION 1116A company has an application that runs on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on Amazon EC2 instances. The application has a UI that uses Amazon DynamoDB and data services that use Amazon S3 as part of the application deployment. The company must ensure that the EKS Pods for the UI can access only Amazon DynamoDB and that the EKS Pods for the data services can access only Amazon S3. The company uses AWS Identity and Access Management (IAM). Which solution meets these requirements?A. Create separate IAM policies for Amazon S3 and DynamoDB access with the required permissions. Attach both IAM policies to the EC2 instance profile. Use role-based access control (RBAC) to control access to Amazon S3 or DynamoDB for the respective EKS Pods.B. Create separate IAM policies for Amazon S3 and DynamoDB access with the required permissions. Attach the Amazon S3 IAM policy directly to the EKS Pods for the data services and the DynamoDB policy to the EKS Pods for the UI.C. Create separate Kubernetes service accounts for the UI and data services to assume an IAM role. Attach the AmazonS3FullAccess policy to the data services account and the AmazonDynamoDBFullAccess policy to the UI service account.D. Create separate Kubernetes service accounts for the UI and data services to assume an IAM role. Use IAM Role for Service Accounts (IRSA) to provide access to the EKS Pods for the UI to Amazon S3 and the EKS Pods for the data services to DynamoDB.Answer: CQUESTION 1117A company needs to give a globally distributed development team secure access to the company's AWS resources in a way that complies with security policies. The company currently uses an on-premises Active Directory for internal authentication. The company uses AWS Organizations to manage multiple AWS accounts that support multiple projects. The company needs a solution to integrate with the existing infrastructure to provide centralized identity management and access control. Which solution will meet these requirements with the LEAST operational overhead?A. Set up AWS Directory Service to create an AWS managed Microsoft Active Directory on AWS. Establish a trust relationship with the on-premises Active Directory. Use IAM roles that are assigned to Active Directory groups to access AWS resources within the company's AWS accounts.B. Create an IAM user for each developer. Manually

manage permissions for each IAM user based on each user's involvement with each project. Enforce multi-factor authentication (MFA) as an additional layer of security.C. Use AD Connector in AWS Directory Service to connect to the on-premises Active Directory. Integrate AD Connector with AWS IAM Identity Center. Configure permissions sets to give each AD group access to specific AWS accounts and resources.D. Use Amazon Cognito to deploy an identity federation solution. Integrate the identity federation solution with the on-premises Active Directory. Use Amazon Cognito to provide access tokens for developers to access AWS accounts and resources.

Answer: CQUESTION 1118A company is developing an application in the AWS Cloud. The application's HTTP API contains critical information that is published in Amazon API Gateway. The critical information must be accessible from only a limited set of trusted IP addresses that belong to the company's internal network. Which solution will meet these requirements?

A. Set up an API Gateway private integration to restrict access to a predefined set of IP addresses.B. Create a resource policy for the API that denies access to any IP address that is not specifically allowed.C. Directly deploy the API in a private subnet. Create a network ACL. Set up rules to allow the traffic from specific IP addresses.D. Modify the security group that is attached to API Gateway to allow inbound traffic from only the trusted IP addresses.

Answer: BQUESTION 1119A company wants to implement new security compliance requirements for its development team to limit the use of approved Amazon Machine Images (AMIs). The company wants to provide access to only the approved operating system and software for all its Amazon EC2 instances. The company wants the solution to have the least amount of lead time for launching EC2 instances. Which solution will meet these requirements?

A. Create a portfolio by using AWS Service Catalog that includes only EC2 instances launched with approved AMIs. Ensure that all required software is preinstalled on the AMIs. Create the necessary permissions for developers to use the portfolio.B. Create an AMI that contains the approved operating system and software by using EC2 Image Builder. Give developers access to that AMI to launch the EC2 instances.C. Create an AMI that contains the approved operating system. Tell the developers to use the approved AMI. Create an Amazon EventBridge rule to run an AWS Systems Manager script when a new EC2 instance is launched. Configure the script to install the required software from a repository.D. Create an AWS Config rule to detect the launch of EC2 instances with an AMI that is not approved. Associate a remediation rule to terminate those instances and launch the instances again with the approved AMI. Use AWS Systems Manager to automatically install the approved software on the launch of an EC2 instance.

Answer: AExplanation: AWS Service Catalog is designed to allow organizations to manage a catalog of approved products (including AMIs) that users can deploy. By creating a portfolio that contains only EC2 instances launched with preapproved AMIs, the company can enforce compliance with the approved operating system and software for all EC2 instances. Service Catalog also streamlines the process of launching EC2 instances, reducing the lead time while ensuring that developers use only the approved configurations.

QUESTION 1120A company needs a solution to automate email ingestion. The company needs to automatically parse email messages, look for email attachments, and save any attachments to an Amazon S3 bucket in near real time. Which solution will meet these requirements?

A. Set up email receiving in Amazon Simple Email Service (Amazon SES). Create a rule set and a receipt rule. Create an AWS Lambda function that Amazon SES can invoke to process the email bodies and attachments.B. Set up email content filtering in Amazon Simple Email Service (Amazon SES). Create a content filtering rule based on sender, recipient, message body, and attachments.C. Set up email receiving in Amazon Simple Email Service (Amazon SES). Configure Amazon SES and S3 Event Notifications to process the email bodies and attachments.D. Create an AWS Lambda function to process the email bodies and attachments. Use Amazon EventBridge to invoke the Lambda function. Configure an EventBridge rule to listen for incoming emails.

Answer: AExplanation: Amazon SES (Simple Email Service) allows for the automatic ingestion of incoming emails. By setting up email receiving in SES and creating a rule set with a receipt rule, you can configure SES to invoke an AWS Lambda function whenever an email is received. The Lambda function can then process the email body and attachments, saving any attachments to an Amazon S3 bucket. This solution is highly scalable, cost-effective, and provides near real-time processing of emails with minimal operational overhead.

QUESTION 1121A logistics company is creating a data exchange platform to share shipment status information with shippers. The logistics company can see all shipment information and metadata. The company distributes shipment data updates to shippers. Each shipper should see only shipment updates that are relevant to their company. Shippers should not see the full detail that is visible to the logistics company. The company creates an Amazon Simple Notification Service (Amazon SNS) topic for each shipper to share data. Some shippers use a mobile app to submit shipment status updates. The company needs to create a data exchange platform that provides each shipper specific access to the data that is relevant to their company. Which solution will meet these requirements with the LEAST operational overhead?

A. Ingest the shipment updates from the mobile app into Amazon Simple Queue Service (Amazon SQS). Publish the updates to the SNS topic. Apply a filter policy to rewrite the body of each message.B. Ingest the shipment updates from the mobile app into Amazon Simple Queue Service (Amazon SQS). Use an AWS Lambda function to consume the updates from Amazon SQS and rewrite the body of each message. Publish the updates to the SNS topic.C. Ingest the shipment updates from the mobile app into a second SNS topic.

Publish the updates to the shipper SNS topic. Apply a filter policy to rewrite the body of each message.D. Ingest the shipment updates from the mobile app into Amazon Simple Queue Service (Amazon SQS). Filter and rewrite the messages in Amazon EventBridge Pipes. Publish the updates to the SNS topic.

Answer: BExplanation: The best solution is to use Amazon SQS to receive updates from the mobile app and process them with an AWS Lambda function. The Lambda function can rewrite the message body as necessary for each shipper and then publish the updates to the appropriate SNS topic for distribution. This setup ensures that each shipper receives only the relevant data and minimizes operational overhead by using managed services.

QUESTION 1122 A company needs to set up a centralized solution to audit API calls to AWS for workloads that run on AWS services and non AWS services. The company must store logs of the audits for 7 years. Which solution will meet these requirements with the LEAST operational overhead?

A. Set up a data lake in Amazon S3. Incorporate AWS CloudTrail logs and logs from non AWS services into the data lake. Use CloudTrail to store the logs for 7 years.

B. Configure custom integrations for AWS CloudTrail Lake to collect and store CloudTrail events from AWS services and non AWS services. Use CloudTrail to store the logs for 7 years.

C. Enable AWS CloudTrail for AWS services. Ingest non AWS services into CloudTrail to store the logs for 7 years.

D. Create new Amazon CloudWatch Logs groups. Send the audit data from non AWS services to the CloudWatch Logs groups. Enable AWS CloudTrail for workloads that run on AWS. Use CloudTrail to store the logs for 7 years.

Answer: BExplanation: AWS CloudTrail Lake is a fully managed service that allows the collection, storage, and querying of CloudTrail events for both AWS and non-AWS services. CloudTrail Lake can be customized to collect logs from various sources, ensuring a centralized audit solution. It also supports long-term storage, so logs can be retained for 7 years, meeting the compliance requirement.

QUESTION 1123 A company needs to migrate a MySQL database from an on-premises data center to AWS within 2 weeks. The database is 180 TB in size. The company cannot partition the database. The company wants to minimize downtime during the migration. The company's internet connection speed is 100 Mbps. Which solution will meet these requirements?

A. Order an AWS Snowball Edge Storage Optimized device. Use AWS Database Migration Service (AWS DMS) and the AWS Schema Conversion Tool (AWS SCT) to migrate the database to Amazon RDS for MySQL and replicate ongoing changes. Send the Snowball Edge device back to AWS to finish the migration. Continue to replicate ongoing changes.

B. Establish an AWS Site-to-Site VPN connection between the data center and AWS. Use AWS Database Migration Service (AWS DMS) and the AWS Schema Conversion Tool (AWS SCT) to migrate the database to Amazon RDS for MySQL and replicate ongoing changes.

C. Establish a 10 Gbps dedicated AWS Direct Connect connection between the data center and AWS. Use AWS DataSync to replicate the database to Amazon S3. Create a script to import the data from Amazon S3 to a new Amazon RDS for MySQL database instance.

D. Use the company's existing internet connection. Use AWS DataSync to replicate the database to Amazon S3. Create a script to import the data from Amazon S3 to a new Amazon RDS for MySQL database instance.

Answer: AExplanation: Given the large size (180 TB) of the database and the time constraint, AWS Snowball Edge Storage Optimized is the best solution. Snowball Edge allows for the physical transfer of large datasets to AWS efficiently without relying on slow internet connections. AWS DMS and SCT can be used to perform ongoing replication of any changes made during the migration, ensuring minimal downtime.

QUESTION 1124 A company is deploying a new gaming application on Amazon EC2 instances. The gaming application needs to have access to shared storage. The company requires a high-performance solution to give the application the ability to use an existing custom protocol to access shared storage. The solution must ensure low latency and must be operationally efficient. Which solution will meet these requirements?

A. Create an Amazon FSx File Gateway. Create a file share that uses the existing custom protocol. Connect the EC2 instances that host the application to the file share.

B. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance. Connect the EC2 instances that host the application to the file share.

C. Create an Amazon Elastic File System (Amazon EFS) file system. Configure the file system to support Lustre. Connect the EC2 instances that host the application to the file system.

D. Create an Amazon FSx for Lustre file system. Connect the EC2 instances that host the application to the file system.

Answer: DExplanation: Amazon FSx for Lustre is a high-performance, fully managed file system that is ideal for applications requiring low-latency access to shared storage, especially in use cases like gaming where high throughput and low latency are essential. It integrates easily with EC2 instances, providing fast and scalable shared storage, and supports custom protocols for specific application needs.

QUESTION 1125 A company is developing a rating system for its ecommerce web application. The company needs a solution to save ratings that users submit in an Amazon DynamoDB table. The company wants to ensure that developers do not need to interact directly with the DynamoDB table. The solution must be scalable and reusable. Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Application Load Balancer (ALB). Create an AWS Lambda function, and set the function as a target group in the ALB. Invoke the Lambda function by using the put\_item method through the ALB.

B. Create an AWS Lambda function. Configure the Lambda function to interact with the DynamoDB table by using the put-item method from Boto3. Invoke the Lambda function from the web application.

C. Create an Amazon Simple Queue Service

(Amazon SQS) queue and an AWS Lambda function that has an SQS trigger type. Instruct the developers to add customer ratings to the SQS queue as JSON messages. Configure the Lambda function to fetch the ratings from the queue and store the ratings in DynamoDB.D. Create an Amazon API Gateway REST API Define a resource and create a new POST method Choose AWS as the integration type, and select DynamoDB as the service. Set the action to PutItem.Answer: DExplanation:Amazon API Gatewayprovides a scalable and reusable solution for interacting with DynamoDB without requiring direct access by developers. By setting up a REST API with a POST methodthat integrates with DynamoDB'sPutItemaction, developers can submit data (such as user ratings) to the DynamoDB table through API Gateway, without having to directly interact with the database. This solution is serverless and minimizes operational overhead.QUESTION 1126A company wants to publish a private website for its on-premises employees. The website consists of several HTML pages and image files. The website must be available only through HTTPS and must be available only to on-premises employees. A solutions architect plans to store the website files in an Amazon S3 bucket. Which solution will meet these requirements?A. Create an S3 bucket policy to deny access when the source IP address is not the public IP address of the on-premises environment Set up an Amazon Route 53 alias record to point to the S3 bucket. Provide the alias record to the on-premises employees to grant the employees access to the website.B. Create an S3 access point to provide website access. Attach an access point policy to deny access when the source IP address is not the public IP address of the on-premises environment. Provide the S3 access point alias to the on-premises employees to grant the employees access to the website.C. Create an Amazon CloudFront distribution that includes an origin access control (OAC) that is configured for the S3 bucket. Use AWS Certificate Manager for SSL. Use AWS WAF with an IP set rule that allows access for the on-premises IP address. Set up an Amazon Route 53 alias record to point to the CloudFront distribution.D. Create an Amazon CloudFront distribution that includes an origin access control (OAC) that is configured for the S3 bucket. Create a CloudFront signed URL for the objects in the bucket. Set up an Amazon Route 53 alias record to point to the CloudFront distribution. Provide the signed URL to the on-premises employees to grant the employees access to the website.Answer: CExplanation:This solution usesCloudFrontto serve the website securely over HTTPS usingAWS Certificate Manager (ACM)for SSL certificates.Origin Access Control (OAC)ensures that only CloudFront can access the S3 bucket directly.AWS WAFwith an IP set rule restricts access to the website, allowing only the on- premises IP address.Route 53is used to create an alias record pointing to the CloudFront distribution. This setup ensures secure, private access to the website with low administrative overhead.QUESTION 1127A manufacturing company runs an order processing application in its VPC. The company wants to securely send messages from the application to an external Salesforce system that uses Open Authorization (OAuth).A solutions architect needs to integrate the company's order processing application with the external Salesforce system.Which solution will meet these requirements?A. Create an Amazon Simple Notification Service (Amazon SNS) topic in a fanout configuration that pushes data to an HTTPS endpoint. Configure the order processing application to publish messages to the SNS topic.B. Create an Amazon Simple Notification Service (Amazon SNS) topic in a fanout configuration that pushes data to an Amazon Data Firehose delivery stream that has a HTTP destination. Configure the order processing application to publish messages to the SNS topic.C. Create an Amazon EventBridge rule and configure an Amazon EventBridge API destination partner Configure the order processing application to publish messages to Amazon EventBridge.D. Create an Amazon Managed Streaming for Apache Kafka (Amazon MSK) topic that has an outbound MSK Connect connector. Configure the order processing application to publish messages to the MSK topic.Answer: CExplanation: AmazonEventBridgeAPI destinations allow you to send data from AWS to external systems, like Salesforce, using HTTP APIs, including those secured with OAuth. This provides a secure and scalable solution for sending messages from the order processing application to Salesforce.QUESTION 1128A company uses an Amazon EC2 Auto Scaling group to host an API. The EC2 instances are in a target group that is associated with an Application Load Balancer (ALB). The company stores data in an Amazon Aurora PostgreSQL database.The API has a weekly maintenance window. The company must ensure that the API returns a static maintenance response during the weekly maintenance window. Which solution will meet this requirement with the LEAST operational overhead?A. Create a table in Aurora PostgreSQL that has fields to contain keys and values. Create a key for a maintenance flag. Set the flag when the maintenance window starts. Configure the API to query the table for the maintenance flag and to return a maintenance response if the flag is set.Reset the flag when the maintenance window is finished.B. Create an Amazon Simple Queue Service (Amazon SQS) queue. Subscribe the EC2 instances to the queue. Publish a message to the queue when the maintenance window starts. Configure the API to return a maintenance message if the instances receive a maintenance start message from the queue.Publish another message to the queue when the maintenance window is finished to restore normal operation.C. Create a listener rule on the ALB to return a maintenance response when the path on a request matches a wildcard. Set the rule priority to one. Perform the maintenance. When the maintenance window is finished, delete the listener rule.D. Create an Amazon Simple Notification Service (Amazon SNS) topic Subscribe the EC2 instances to the topic Publish a message to the

topic when the maintenance window starts. Configure the API to return a maintenance response if the instances receive the maintenance start message from the topic. Publish another message to the topic when the maintenance window finishes to restore normal operation.

Answer: CExplanation: Creating a listener rule on the Application Load Balancer (ALB) to return a maintenance response during the maintenance window is the most straightforward solution with the least operational overhead. The rule can be configured to match all incoming requests and return a custom response, and it can be easily removed once maintenance is complete.

QUESTION 1129 An online education platform experiences lag and buffering during peak usage hours, when thousands of students access video lessons concurrently. A solutions architect needs to improve the performance of the education platform. The platform needs to handle unpredictable traffic surges without losing responsiveness. The platform must provide smooth video playback performance at all times. The platform must create multiple copies of each video lesson and store the copies in various bitrates to serve users who have different internet speeds. The smallest video size is 7 GB. Which solution will meet these requirements MOST cost-effectively?

- A. Use Amazon ElastiCache to cache videos in all the required bitrates. Use AWS Lambda functions to process the videos and to convert the videos to the required bitrates.
- B. Create an Auto Scaling group that includes Amazon EC2 instances that are sized to meet peak loads. Use the Auto Scaling group to serve videos. Use the Auto Scaling group to convert the videos to the required bitrates.
- C. Store a copy of every video in every required bitrate in an Amazon S3 bucket. Use a single Amazon EC2 instance to serve the videos.
- D. Use Amazon Kinesis Video Streams to store and serve the videos. Use AWS Lambda functions to process the videos and to convert the videos to the required bitrates.

Answer: CExplanation: The most cost-effective solution for serving video content with different bitrates is to store multiple versions of each video in Amazon S3. S3 provides scalable and cost-effective storage for large media files. Serving the videos from a single Amazon EC2 instance ensures low-latency delivery, and S3 storage helps minimize costs.

QUESTION 1130 A company has Amazon EC2 instances in multiple AWS Regions. The instances all store and retrieve confidential data from the same Amazon S3 bucket. The company wants to improve the security of its current architecture. The company wants to ensure that only the Amazon EC2 instances within its VPC can access the S3 bucket. The company must block all other access to the bucket. Which solution will meet this requirement?

- A. Use IAM policies to restrict access to the S3 bucket.
- B. Use server-side encryption (SSE) to encrypt data in the S3 bucket at rest. Store the encryption key on the EC2 instances.
- C. Create a VPC endpoint for Amazon S3. Configure an S3 bucket policy to allow connections only from the endpoint.
- D. Use AWS Key Management Service (AWS KMS) with customer-managed keys to encrypt the data before sending the data to the S3 bucket.

Answer: CExplanation: Creating a VPC endpoint for S3 and configuring a bucket policy to allow access only from the endpoint ensures that only EC2 instances within the VPC can access the S3 bucket. This solution improves security by restricting access at the network level without the need for public internet access.

QUESTION 1131 A company recently launched a new product that is highly available in one AWS Region. The product consists of an application that runs on Amazon Elastic Container Service (Amazon ECS), a public Application Load Balancer (ALB), and an Amazon DynamoDB table. The company wants a solution that will make the application highly available across Regions. Which combination of steps will meet these requirements? (Choose three.)

- A. In a different Region, deploy the application to a new ECS cluster that is accessible through a new ALB.
- B. Create an Amazon Route 53 failover record.
- C. Modify the DynamoDB table to create a DynamoDB global table.
- D. In the same Region, deploy the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is accessible through a new ALB.
- E. Modify the DynamoDB table to create global secondary indexes (GSIs).

F. Create an AWS PrivateLink endpoint for the application.

Answer: ABCExplanation: To make the application highly available across regions, Deploy the application in a different region using a new ECS cluster and ALB to ensure regional redundancy. Use Route 53 failover routing to automatically direct traffic to the healthy region in case of failure. Use DynamoDB Global Tables to ensure the database is replicated and available across multiple regions, supporting read and write operations in each region.

QUESTION 1132 A company runs a payment processing system in the AWS Cloud. Sometimes when a payment fails because of insufficient funds or technical issues, users attempt to resubmit the payment. Sometimes payment resubmissions invoke multiple payment messages for the same payment ID. A solutions architect needs to ensure that the payment processing system receives payment messages that have the same payment ID sequentially, according to when the messages were generated. The processing system must process the messages in the order in which the messages are received. The solution must retain all payment messages for 10 days for analytics. Which solutions will meet these requirements? (Choose two.)

- A. Write the payment messages to an Amazon DynamoDB table that uses the payment ID as the partition key.
- B. Write the payment messages to an Amazon Kinesis data stream that uses the payment ID as the partition key.
- C. Write the payment messages to an Amazon ElastiCache for Memcached cluster that uses the payment ID as the key.
- D. Write the payment messages to an Amazon Simple Queue Service (Amazon SQS) queue. Set the message attribute to use the payment ID.
- E. Write the payment messages to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Set the message group to use the payment ID.

Answer: BEExplanation: Both Amazon Kinesis and SQS FIFO queues ensure the sequential

processing of messages. By using the payment ID as the partition key in Kinesis or as the message group in the SQS FIFO queue, messages are processed in order. Both solutions also allow for long-term retention (up to 10 days) of messages, making them suitable for this payment processing use case.

**QUESTION 1133**A company stores customer data in a multitenant Amazon S3 bucket. Each customer's data is stored in a prefix that is unique to the customer. The company needs to migrate data for specific customers to a new, dedicated S3 bucket that is in the same AWS Region as the source bucket. The company must preserve object metadata such as creation date and version IDs. After the migration is finished, the company must delete the source data for the migrated customers from the original multitenant S3 bucket. Which combination of solutions will meet these requirements with the LEAST overhead? (Choose three.)

- A. Create a new S3 bucket as a destination bucket. Enable versioning on the new bucket.
- B. Use S3 batch operations to copy objects from the specified prefixes to the destination bucket.
- C. Use the S3 CopyObject API, and create a script to copy data to the destination S3 bucket.
- D. Configure S3 Same-Region Replication (SRR) to replicate existing data from the specified prefixes in the source bucket to the destination bucket.
- E. Configure AWS DataSync to migrate data from the specified prefixes in the source bucket to the destination bucket.
- F. Use an S3 Lifecycle policy to delete objects from the source bucket after the data is migrated to the destination bucket.

**Answer:** ABFExplanation: The combination of these solutions provides an efficient and automated way to migrate data while preserving metadata and ensuring cleanup: Create a new S3 bucket with versioning enabled (Option A) to preserve object metadata like version IDs during migration. Use S3 batch operations (Option B) to efficiently copy data from specific prefixes in the source bucket to the destination bucket, ensuring minimal overhead. Use an S3 Lifecycle policy (Option F) to automatically delete the data from the source bucket after it has been migrated, reducing manual intervention.

**QUESTION 1134**A media company is using video conversion tools that run on Amazon EC2 instances. The video conversion tools run on a combination of Windows EC2 instances and Linux EC2 instances. Each video file is tens of gigabytes in size. The video conversion tools must process the video files in the shortest possible amount of time. The company needs a single, centralized file storage solution that can be mounted on all the EC2 instances that host the video conversion tools. Which solution will meet these requirements?

- A. Deploy Amazon FSx for Windows File Server with hard disk drive (HDD) storage.
- B. Deploy Amazon FSx for Windows File Server with solid state drive (SSD) storage.
- C. Deploy Amazon Elastic File System (Amazon EFS) with Max I/O performance mode.
- D. Deploy Amazon Elastic File System (Amazon EFS) with General Purpose performance mode.

**Answer:** CExplanation: Amazon EFS with Max I/O performance mode is designed for workloads that require high levels of parallelism, such as video processing across multiple EC2 instances. EFS provides shared file storage that can be mounted on both Windows and Linux EC2 instances, and the Max I/O mode ensures the best performance for handling large files and concurrent access across multiple instances.

**QUESTION 1135**A company is designing a microservice-based architecture for a new application on AWS. Each microservice will run on its own set of Amazon EC2 instances. Each microservice will need to interact with multiple AWS services such as Amazon S3 and Amazon Simple Queue Service (Amazon SQS). The company wants to manage permissions for each EC2 instance based on the principle of least privilege. Which solution will meet this requirement?

- A. Assign an IAM user to each micro-service. Use access keys stored within the application code to authenticate AWS service requests.
- B. Create a single IAM role that has permission to access all AWS services. Associate the IAM role with all EC2 instances that run the microservices.
- C. Use AWS Organizations to create a separate account for each microservice. Manage permissions at the account level.
- D. Create individual IAM roles based on the specific needs of each microservice. Associate the IAM roles with the appropriate EC2 instances.

**Answer:** DExplanation: When designing a microservice architecture where each microservice interacts with different AWS services, it's essential to follow the principle of least privilege. This means granting each microservice only the permissions it needs to perform its tasks, reducing the risk of unauthorized access or accidental actions. The recommended approach is to create individual IAM roles with policies that grant each microservice the specific permissions it requires. Then, these roles should be associated with the EC2 instances that run the corresponding microservice. By doing so, each EC2 instance will assume its specific IAM role, and permissions will be automatically managed by AWS. IAM roles provide temporary credentials via the instance metadata service, eliminating the need to hard-code credentials in your application code, which enhances security.

**QUESTION 1136**A media company hosts its video processing workload on AWS. The workload uses Amazon EC2 instances in an Auto Scaling group to handle varying levels of demand. The workload stores the original videos and the processed videos in an Amazon S3 bucket. The company wants to ensure that the video processing workload is scalable. The company wants to prevent failed processing attempts because of resource constraints. The architecture must be able to handle sudden spikes in video uploads without impacting the processing capability. Which solution will meet these requirements with the LEAST overhead?

- A. Migrate the workload from Amazon EC2 instances to AWS Lambda functions. Configure an Amazon S3 event notification to invoke the Lambda functions when a new video is uploaded. Configure the Lambda functions to process videos directly and to save processed videos back to the S3 bucket.
- B. Migrate the workload from Amazon EC2 instances to AWS Lambda functions. Use Amazon S3 to

invoke an Amazon Simple Notification Service (Amazon SNS) topic when a new video is uploaded. Subscribe the Lambda functions to the SNS topic. Configure the Lambda functions to process the videos asynchronously and to save processed videos back to the S3 bucket.C. Configure an Amazon S3 event notification to send a message to an Amazon Simple Queue Service (Amazon SQS) queue when a new video is uploaded. Configure the existing Auto Scaling group to poll the SQS queue, process the videos, and save processed videos back to the S3 bucket.D. Configure an Amazon S3 upload trigger to invoke an AWS Step Functions state machine when a new video is uploaded. Configure the state machine to orchestrate the video processing workflow by placing a job message in the Amazon SQS queue. Configure the job message to invoke the EC2 instances to process the videos. Save processed videos back to the S3 bucket.

**Answer:** CExplanation: This solution addresses the scalability needs of the workload while preventing failed processing attempts due to resource constraints. Amazon S3 event notifications can be used to trigger a message to an SQS queue whenever a new video is uploaded. The existing Auto Scaling group of EC2 instances can poll the SQS queue, ensuring that the EC2 instances only process videos when there is a job in the queue. SQS decouples the video upload and processing steps, allowing the system to handle sudden spikes in video uploads without overloading EC2 instances. The use of Auto Scaling ensures that the EC2 instances can scale in or out based on the demand, maintaining cost efficiency while avoiding processing failures due to insufficient resources.

**QUESTION 1137** A company uses a set of Amazon EC2 instances to host a website. The website uses an Amazon S3 bucket to store images and media files. The company wants to automate website infrastructure creation to deploy the website to multiple AWS Regions. The company also wants to provide the EC2 instances access to the S3 bucket so the instances can store and access data by using AWS Identity and Access Management (IAM). Which solution will meet these requirements MOST securely?

A. Create an AWS CloudFormation template for the web server EC2 instances. Save an IAM access key in the UserData section of the AWS::EC2::Instance entity in the CloudFormation template.

B. Create a file that contains an IAM secret access key and access key ID. Store the file in a new S3 bucket. Create an AWS CloudFormation template. In the template, create a parameter to specify the location of the S3 object that contains the access key and access key ID.

C. Create an IAM role and an IAM access policy that allows the web server EC2 instances to access the S3 bucket. Create an AWS CloudFormation template for the web server EC2 instances that contains an IAM instance profile entity that references the IAM role and the IAM access policy.

D. Create a script that retrieves an IAM secret access key and access key ID from IAM and stores them on the web server EC2 instances. Include the script in the UserData section of the AWS::EC2::Instance entity in an AWS CloudFormation template.

**Answer:** CExplanation: The most secure solution for allowing EC2 instances to access an S3 bucket is by using IAM roles. An IAM role can be created with an access policy that grants the required permissions (e.g., to read and write to the S3 bucket). The IAM role is then associated with the EC2 instances through an IAM instance profile. By associating the role with the instances, the EC2 instances can securely assume the role and receive temporary credentials via the instance metadata service. This avoids the need to store credentials (such as access keys) on the instances or within the application, enhancing security and reducing the risk of credentials being exposed. AWS CloudFormation can be used to automate the creation of the entire infrastructure, including EC2 instances, IAM roles, and associated policies.

**QUESTION 1138** A finance company is migrating its trading platform to AWS. The trading platform processes a high volume of market data and processes stock trades. The company needs to establish a consistent, low-latency network connection from its on-premises data center to AWS. The company will host resources in a VPC. The solution must not use the public internet. Which solution will meet these requirements?

A. Use AWS Client VPN to connect the on-premises data center to AWS.

B. Use AWS Direct Connect to set up a connection from the on-premises data center to AWS.

C. Use AWS PrivateLink to set up a connection from the on-premises data center to AWS.

D. Use AWS Site-to-Site VPN to connect the on-premises data center to AWS.

**Answer:** BExplanation: AWS Direct Connect is the best solution for establishing a consistent, low-latency connection from an on-premises data center to AWS without using the public internet. Direct Connect offers dedicated, high-throughput, and low-latency network connections, which are ideal for performance-sensitive applications like a trading platform that processes high volumes of market data and stock trades. Direct Connect provides a private connection to your AWS VPC, ensuring that data doesn't traverse the public internet, which enhances both security and performance consistency.

**QUESTION 1139** A solutions architect is designing the architecture for a company website that is composed of static content. The company's target customers are located in the United States and Europe. Which architecture should the solutions architect recommend to MINIMIZE cost?

A. Store the website files on Amazon S3 in the us-east-2 Region. Use an Amazon CloudFront distribution with the price class configured to limit the edge locations in use.

B. Store the website files on Amazon S3 in the us-east-2 Region. Use an Amazon CloudFront distribution with the price class configured to maximize the use of edge locations.

C. Store the website files on Amazon S3 in the us-east-2 Region and the eu-west-1 Region. Use an Amazon CloudFront geolocation routing policy to route requests to the closest Region to the user.

D. Store the website files on Amazon S3 in the us-east-2 Region and the eu-west-1 Region. Use an Amazon CloudFront distribution with an Amazon Route 53 latency routing policy to route requests to the closest

Region to the user. Answer: AExplanation: The question focuses on minimizing costs while serving static content to users in the US and Europe. Option A uses a single S3 bucket and configures CloudFront to limit edge locations, reducing costs by using fewer edge locations while still improving performance. QUESTION 1140 A company hosts a database that runs on an Amazon RDS instance deployed to multiple Availability Zones. A periodic script negatively affects a critical application by querying the database. How can application performance be improved with minimal costs? A. Add functionality to the script to identify the instance with the fewest active connections and query that instance. B. Create a read replica of the database. Configure the script to query only the read replica. C. Instruct the development team to manually export new entries at the end of the day. D. Use Amazon ElastiCache to cache the common queries the script runs. Answer: BQUESTION 1141 A company has developed an API using Amazon API Gateway REST API and AWS Lambda. How can latency be reduced for users worldwide? A. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Enable content encoding to compress data in transit. B. Deploy the REST API as a Regional API endpoint. Enable caching. Enable content encoding to compress data in transit. C. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Configure reserved concurrency for Lambda functions. D. Deploy the REST API as a Regional API endpoint. Enable caching. Configure reserved concurrency for Lambda functions. Answer: AQUESTION 1142 How can a law firm make files publicly readable while preventing modifications or deletions until a specific future date? A. Upload files to an Amazon S3 bucket configured for static website hosting. Grant read-only IAM permissions to any AWS principals. B. Create an S3 bucket. Enable S3 Versioning. Use S3 Object Lock with a retention period. Create a CloudFront distribution. Use a bucket policy to restrict access. C. Create an S3 bucket. Enable S3 Versioning. Configure an event trigger with AWS Lambda to restore modified objects from a private S3 bucket. D. Upload files to an S3 bucket for static website hosting. Use S3 Object Lock with a retention period. Grant read-only IAM permissions. Answer: BQUESTION 1143 A media company hosts a web application on AWS for uploading videos. Only authenticated users should upload within a specified time frame after authentication. Which solution will meet these requirements with the LEAST operational overhead? A. Configure the application to generate IAM temporary security credentials for authenticated users. B. Create an AWS Lambda function that generates pre-signed URLs when a user authenticates. C. Develop a custom authentication service that integrates with Amazon Cognito to control and log direct S3 bucket access through the application. D. Use AWS Security Token Service (AWS STS) to assume a pre-defined IAM role that grants authenticated users temporary permissions to upload videos directly to the S3 bucket. Answer: BQUESTION 1144 A company needs to ingest and analyze telemetry data from vehicles at scale for machine learning and reporting. Which solution will meet these requirements? A. Use Amazon Timestream for LiveAnalytics to store data points. Grant Amazon SageMaker permission to access the data. Use Amazon QuickSight to visualize the data. B. Use Amazon DynamoDB to store data points. Use DynamoDB Connector to ingest data into Amazon EMR for processing. Use Amazon QuickSight to visualize the data. C. Use Amazon Neptune to store data points. Use Amazon Kinesis Data Streams to ingest data into a Lambda function for processing. Use Amazon QuickSight to visualize the data. D. Use Amazon Timestream for LiveAnalytics to store data points. Grant Amazon SageMaker permission to access the data. Use Amazon Athena to visualize the data. Answer: AQUESTION 1145 A company runs an application on EC2 instances that need access to RDS credentials stored in AWS Secrets Manager. Which solution meets this requirement? A. Create an IAM role, and attach the role to each EC2 instance profile. Use an identity-based policy to grant the role access to the secret. B. Create an IAM user, and attach the user to each EC2 instance profile. Use a resource-based policy to grant the user access to the secret. C. Create a resource-based policy for the secret. Use EC2 Instance Connect to access the secret. D. Create an identity-based policy for the secret. Grant direct access to the EC2 instances. Answer: AQUESTION 1146 A company needs a cloud-based solution for backup, recovery, and archiving while retaining encryption key material control. Which combination of solutions will meet these requirements? (Select TWO) A. Create an AWS Key Management Service (AWS KMS) key without key material. Import the company's key material into the KMS key. B. Create an AWS KMS encryption key that contains key material generated by AWS KMS. C. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Use S3 Bucket Keys with AWS KMS keys. D. Store the data in an Amazon S3 Glacier storage class. Use server-side encryption with customer-provided keys (SSE-C). E. Store the data in AWS Snowball devices. Use server-side encryption with AWS KMS keys (SSE-KMS). Answer: ADQUESTION 1147 A website uses EC2 instances with Auto Scaling and EFS. How can the company optimize costs? A. Reconfigure the Auto Scaling group to set a desired number of instances. Turn off scheduled scaling. B. Create a new launch template version that uses larger EC2 instances. C. Reconfigure the Auto Scaling group to use a target tracking scaling policy. D. Replace the EFS volume with instance store volumes. Answer: CQUESTION 1148 A company is developing a social media application that must scale to meet demand spikes and handle ordered processes. Which AWS services meet these requirements? A. ECS with Fargate, RDS, and SQS for decoupling. B. ECS with Fargate, RDS, and SNS for decoupling. C. DynamoDB, Lambda, DynamoDB Streams, and Step Functions. D. Elastic Beanstalk, RDS, and SNS for decoupling. Answer: AQUESTION 1149 How

can a company detect and notify security teams about PII in S3 buckets?A. Use Amazon Macie. Create an EventBridge rule for SensitiveData findings and send an SNS notification.B. Use Amazon GuardDuty. Create an EventBridge rule for CRITICAL findings and send an SNS notification.C. Use Amazon Macie. Create an EventBridge rule for SensitiveData:S3Object/Personal findings and send an SQS notification.D. Use Amazon GuardDuty. Create an EventBridge rule for CRITICAL findings and send an SQS notification.

Answer: A

QUESTION 1150A company runs HPC workloads requiring high IOPS.Which combination of steps will meet these requirements? (Select TWO)

- A. Use Amazon EFS as a high-performance file system.
- B. Use Amazon FSx for Lustre as a high-performance file system.
- C. Create an Auto Scaling group of EC2 instances. Use Reserved Instances. Configure a spread placement group. Use AWS Batch for analytics.
- D. Use Mountpoint for Amazon S3 as a high-performance file system.
- E. Create an Auto Scaling group of EC2 instances. Use mixed instance types and a cluster placement group. Use Amazon EMR for analytics.

Answer: B

QUESTION 1151How can trade data from DynamoDB be ingested into an S3 data lake for near real-time analysis?A. Use DynamoDB Streams to invoke a Lambda function that writes to S3.B. Use DynamoDB Streams to invoke a Lambda function that writes to Data Firehose, which writes to S3.C. Enable Kinesis Data Streams on DynamoDB. Configure it to invoke a Lambda function that writes to S3.D. Enable Kinesis Data Streams on DynamoDB. Use Data Firehose to write to S3.

Answer: A

QUESTION 1152How can DynamoDB data be made available for long-term analytics with minimal operational overhead?A. Configure DynamoDB incremental exports to S3.B. Configure DynamoDB Streams to write records to S3.C. Configure EMR to copy DynamoDB data to S3.D. Configure EMR to copy DynamoDB data to HDFS.

Answer: A

QUESTION 1153A company runs a Microsoft Windows SMB file share on-premises to support an application. The company wants to migrate the application to AWS. The company wants to share storage across multiple Amazon EC2 instances.Which solutions will meet these requirements with the LEAST operational overhead? (Choose two.)

- A. Create an Amazon Elastic File System (Amazon EFS) file system with elastic throughput.
- B. Create an Amazon FSx for NetApp ONTAP file system.
- C. Use Amazon Elastic Block Store (Amazon EBS) to create a self-managed Windows file share on the instances.
- D. Create an Amazon FSx for Windows File Server file system.
- E. Create an Amazon FSx for OpenZFS file system.

Answer: A

QUESTION 1154A solutions architect needs to implement a solution that can handle up to 5,000 messages per second. The solution must publish messages as events to multiple consumers. The messages are upto 500 KB in size. The message consumers need to have the ability to use multiple programming languages to consume the messages with minimal latency. The solution must retain published messages for more than 3 months. The solution must enforce strict ordering of the messages.Which solution will meet these requirements?A. Publish messages to an Amazon Kinesis Data Streams data stream. Enable enhanced fan-out. Ensure that consumers ingest the data stream by using dedicated throughput.

B. Publish messages to an Amazon Simple Notification Service (Amazon SNS) topic. Ensure that consumers use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to subscribe to the topic.

C. Publish messages to Amazon EventBridge. Allow each consumer to create rules to deliver messages to the consumer's own target.

D. Publish messages to an Amazon Simple Notification Service (Amazon SNS) topic. Ensure that consumers use Amazon Data Firehose to subscribe to the topic.

Answer: A

QUESTION 1155A company is planning to migrate an on-premises online transaction processing (OLTP) database that uses MySQL to an AWS managed database management system. Several reporting and analytics applications use the on-premises database heavily on weekends and at the end of each month. The cloud-based solution must be able to handle read-heavy surges during weekends and at the end of each month.Which solution will meet these requirements?A. Migrate the database to an Amazon Aurora MySQL cluster. Configure Aurora Auto Scaling to use replicas to handle surges.

B. Migrate the database to an Amazon EC2 instance that runs MySQL. Use an EC2 instance type that has ephemeral storage. Attach Amazon EBS Provisioned IOPS SSD (io2) volumes to the instance.

C. Migrate the database to an Amazon RDS for MySQL database. Configure the RDS for MySQL database for a Multi-AZ deployment, and set up auto scaling.

D. Migrate from the database to Amazon Redshift. Use Amazon Redshift as the database for both OLTP and analytics applications.

Answer: A

QUESTION 1156A company runs an order management application on AWS. The application allows customers to place orders and pay with a credit card. The company uses an Amazon CloudFront distribution to deliver the application. A security team has set up logging for all incoming requests. The security team needs a solution to generate an alert if any user modifies the logging configuration. Which combination of solutions will meet these requirements? (Choose two.)

- A. Configure an Amazon EventBridge rule that is invoked when a user creates or modifies a CloudFront distribution. Add the AWS Lambda function as a target of the EventBridge rule.
- B. Create an Application Load Balancer (ALB). Enable AWS WAF rules for the ALB. Configure an AWS Config rule to detect security violations.
- C. Create an AWS Lambda function to detect changes in CloudFront distribution logging. Configure the Lambda function to use Amazon Simple Notification Service (Amazon SNS) to send notifications to the security team.
- D. Set up Amazon GuardDuty. Configure GuardDuty to monitor findings from the CloudFront distribution. Create an AWS Lambda function to address the findings.
- E. Create a private API in Amazon API Gateway. Use AWS WAF rules to protect the private API from

common security problems.

Answer: ACQUESTION 1157A company has a website that handles dynamic traffic loads. The website architecture is based on Amazon EC2 instances in an Auto Scaling group that is configured to use scheduled scaling. Each EC2 instance runs code from an Amazon Elastic File System (Amazon EFS) volume and stores shared data back to the same volume. The company wants to optimize costs for the website. Which solution will meet this requirement?

A. Reconfigure the Auto Scaling group to set a desired number of instances. Turn off scheduled scaling.

B. Create a new launch template version for the Auto Scaling group that uses larger EC2 instances.

C. Reconfigure the Auto Scaling group to use a target tracking scaling policy.

D. Replace the EFS volume with instance store volumes.

Answer: CQUESTION 1158A company wants to provide a third-party system that runs in a private data center with access to its AWS account. The company wants to call AWS APIs directly from the third-party system. The company has an existing process for managing digital certificates. The company does not want to use SAML or OpenID Connect (OIDC) capabilities and does not want to store long-term AWS credentials. Which solution will meet these requirements?

A.

Configure mutual TLS to allow authentication of the client and server sides of the communication channel.

B. Configure AWS Signature Version 4 to authenticate incoming HTTPS requests to AWS APIs.

C. Configure Kerberos to exchange tickets for assertions that can be validated by AWS APIs.

D. Configure AWS Identity and Access Management (IAM) Roles Anywhere to exchange X.509 certificates for AWS credentials to interact with AWS APIs.

Answer: DQUESTION 1159A company is migrating a new application from an on-premises data center to a new VPC in the AWS Cloud. The company has multiple AWS accounts and VPCs that share many subnets and applications. The company wants to have fine-grained access control for the new application. The company wants to ensure that all network resources across accounts and VPCs that are granted permission to access the new application can access the application. Which solution will meet these requirements?

A. Set up a VPC peering connection for each VPC that needs access to the new application VPC. Update route tables in each VPC to enable connectivity.

B. Deploy a transit gateway in the account that hosts the new application. Share the transit gateway with each account that needs to connect to the application. Update route tables in the VPC that hosts the new application and in the transit gateway to enable connectivity.

C. Use an AWS PrivateLink endpoint service to make the new application accessible to other VPCs. Control access to the application by using an endpoint policy.

D. Use an Application Load Balancer (ALB) to expose the new application to the internet. Configure authentication and authorization processes to ensure that only specified VPCs can access the application.

Answer: BQUESTION 1160A company is enhancing the security of its AWS environment, where the company stores a significant amount of sensitive customer data. The company needs a solution that automatically identifies and classifies sensitive data that is stored in multiple Amazon S3 buckets. The solution must automatically respond to data breaches and alert the company's security team through email immediately when noncompliant data is found. Which solution will meet these requirements?

A. Use Amazon GuardDuty.

Configure an AWS Lambda function to route alerts to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team to the SNS topic.

B. Use Amazon GuardDuty. Configure an AWS Lambda function to route alerts to an Amazon Simple Queue Service (Amazon SQS) queue. Configure a second Lambda function to periodically poll the SQS queue and to send emails to the security team by using Amazon Simple Email Service (Amazon SES).

C. Use Amazon Macie. Integrate Amazon EventBridge with Macie, and configure EventBridge to send alerts to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team to the SNS topic.

D. Use Amazon Macie. Integrate Amazon EventBridge with Macie, and configure EventBridge to route alerts to an Amazon Simple Queue Service (Amazon SQS) queue. Configure an AWS Lambda function to periodically poll the SQS queue and to send alerts to the security team by using Amazon Simple Email Service (Amazon SES).

Answer: CQUESTION 1161An ecommerce company is planning to migrate an on-premises Microsoft SQL Server database to the AWS Cloud. The company needs to migrate the database to SQL Server Always On availability groups. The cloud-based solution must be highly available. Which solution will meet these requirements?

A. Deploy three Amazon EC2 instances with SQL Server across three Availability Zones. Attach one Amazon Elastic Block Store (Amazon EBS) volume to the EC2 instances.

B. Migrate the database to Amazon RDS for SQL Server. Configure a Multi-AZ deployment and read replicas.

C. Deploy three Amazon EC2 instances with SQL Server across three Availability Zones. Use Amazon FSx for Windows File Server as the storage tier.

D.

Deploy three Amazon EC2 instances with SQL Server across three Availability Zones. Use Amazon S3 as the storage tier.

Answer: CQUESTION 1162A company is developing a highly available natural language processing (NLP) application. The application handles large volumes of concurrent requests. The application performs NLP tasks such as entity recognition, sentiment analysis, and key phrase extraction on text data. The company needs to store data that the application processes in a highly available and scalable database. Which solution will meet these requirements?

A. Create an Amazon API Gateway REST API endpoint to handle incoming requests. Configure the REST API to invoke an AWS Lambda function for each request. Configure the Lambda function to call Amazon Comprehend to perform NLP tasks on the text data. Store the processed data in Amazon DynamoDB.

B. Create an Amazon API Gateway HTTP API endpoint to handle incoming requests. Configure the HTTP API to invoke an AWS Lambda

function for each request. Configure the Lambda function to call Amazon Translate to perform NLP tasks on the text data. Store the processed data in Amazon ElastiCache.C. Create an Amazon SQS queue to buffer incoming requests. Deploy the NLP application on Amazon EC2 instances in an Auto Scaling group. Use Amazon Comprehend to perform NLP tasks. Store the processed data in an Amazon RDS database.D. Create an Amazon API Gateway WebSocket API endpoint to handle incoming requests.Configure the WebSocket API to invoke an AWS Lambda function for each request. Configure the Lambda function to call Amazon Textract to perform NLP tasks on the text data. Store the processed data in Amazon ElastiCache.Answer: AQUESTION 1163A company has developed an API using an Amazon API Gateway REST API and AWS Lambda functions. The API serves static and dynamic content to users worldwide. The company wants to decrease the latency of transferring content for API requests.Which solution will meet these requirements?A. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Enable content encoding in the API definition to compress the application data in transit.B. Deploy the REST API as a Regional API endpoint. Enable caching. Enable content encoding in the API definition to compress the application data in transit.C. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Configure reserved concurrency for the Lambda functions.D. Deploy the REST API as a Regional API endpoint. Enable caching. Configure reserved concurrency for the Lambda functions.Answer: AQUESTION 1164A gaming company is building an application that uses a database to store user data. The company wants the database to have an active-active configuration that allows data writes to a secondary AWS Region. The database must achieve a sub-second recovery point objective (RPO).Which solution will meet these requirements?A. Deploy an Amazon ElastiCache (Redis OSS) cluster. Configure a global data store for disaster recovery. Configure the ElastiCache cluster to cache data from an Amazon RDS database that is deployed in the primary Region.B. Deploy an Amazon DynamoDB table in the primary Region and the secondary Region. Configure Amazon DynamoDB Streams to invoke an AWS Lambda function to write changes from the table in the primary Region to the table in the secondary Region.C. Deploy an Amazon Aurora MySQL database in the primary Region. Configure a global database for the secondary Region.D. Deploy an Amazon DynamoDB table in the primary Region. Configure global tables for the secondary Region.Answer: DQUESTION 1165A company wants to implement a data lake in the AWS Cloud. The company must ensure that only specific teams have access to sensitive data in the data lake. The company must have row-level access control for the data lake.Which solution meet these requirements?A. Use Amazon RDS to store the data. Use IAM roles and permissions for data governance and access control.B. Use Amazon Redshift to store the data. Use IAM roles and permissions for data governance and access control.C. Use Amazon S3 to store the data. Use AWS Lake Formation for data governance and access control.D. Use AWS Glue Catalog to store the data. Use AWS Glue DataBrew for data governance and access control.Answer: CQUESTION 1166A company hosts a multi-tier inventory reporting application on AWS. The company needs a cost- effective solution to generate inventory reports on demand. Admin users need to have the ability to generate new reports. Reports take approximately 5-10 minutes to finish. The application must send reports to the email address of the admin user who generates each report.Which solution meet these requirements?A. Use Amazon Elastic Container Service (Amazon ECS) to host the report generation code. Use an Amazon API Gateway HTTP API to invoke the code. Use Amazon Simple Email Service (Amazon SES) to send the reports to admin users.B. Use Amazon EventBridge to invoke a scheduled AWS Lambda function to generate the reports.Use Amazon Simple Notification Service (Amazon SNS) to send the reports to admin users.C. Use Amazon Elastic Kubernetes Service (Amazon EKS) to host the report generation code.Use an Amazon API Gateway REST API to invoke the code. Use Amazon Simple Notification Service (Amazon SNS) to send the reports to admin users.D. Create an AWS Lambda function to generate the reports. Use a function URL to invoke the function. Use Amazon Simple Email Service (Amazon SES) to send the reports to admin users.Answer: DQUESTION 1167A company that has multiple AWS accounts maintains an on-premises Microsoft Active Directory. The company needs a solution to implement Single Sign-On for its employees. The company wants to use AWS IAM Identity Center.The solution must meet the following requirements:- Allow users to access AWS accounts and third-party applications by using existing Active Directory credentials.- Enforce multi-factor authentication (MFA) to access AWS accounts.- Centrally manage permissions to access AWS accounts and applications.Which solution meet these requirements?A. Create an IAM identity provider for Active Directory in each AWS account. Ensure that Active Directory users and groups access AWS accounts directly through IAM roles. Use IAM Identity Center to enforce MFA in each account for all users.B. Use AWS Directory Service to create a new AWS Managed Microsoft AD Active Directory.Configure IAM Identity Center in each account to use the new AWS Managed Microsoft AD Active Directory as the identity source. Use IAM Identity Center to enforce MFA for all users.C. Use IAM Identity Center with the existing Active Directory as the identity source. Enforce MFA for all users. Use AWS Organizations and Active Directory groups to manage access permissions for AWS accounts and application access.D. Use AWS Lambda functions to periodically synchronize Active Directory users and groups with IAM users and groups in each AWS account. Use IAM roles and policies to manage application access. Create a second Lambda function to enforce MFA.Answer: CQUESTION

1168A company runs an order management application on AWS. The application allows customers to place orders and pay with a credit card. The company uses an Amazon CloudFront distribution to deliver the application. A security team has set up logging for all incoming requests. The security team needs a solution to generate an alert if any user modifies the logging configuration. Which solution meet these requirements? (Choose two.)  
A. Configure an Amazon EventBridge rule that is invoked when a user creates or modifies a CloudFront distribution. Add the AWS Lambda function as a target of the EventBridge rule.  
B. Create an Application Load Balancer (ALB). Enable AWS WAF rules for the ALB. Configure an AWS Config rule to detect security violations.  
C. Create an AWS Lambda function to detect changes in CloudFront distribution logging. Configure the Lambda function to use Amazon Simple Notification Service (Amazon SNS) to send notifications to the security team.  
D. Set up Amazon GuardDuty. Configure GuardDuty to monitor findings from the CloudFront distribution. Create an AWS Lambda function to address the findings.  
E. Create a private API in Amazon API Gateway. Use AWS WAF rules to protect the private API from common security problems.

Answer: ACQUESTION 1169A company has developed an API by using an Amazon API Gateway REST API and AWS Lambda functions. The API serves static content and dynamic content to users worldwide. The company wants to decrease the latency of transferring the content for API requests. Which solution will meet these requirements?  
A. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Enable content encoding in the API definition to compress the application data in transit.  
B. Deploy the REST API as a Regional API endpoint. Enable caching. Enable content encoding in the API definition to compress the application data in transit.  
C. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Configure reserved concurrency for the Lambda functions.  
D. Deploy the REST API as a Regional API endpoint. Enable caching. Configure reserved concurrency for the Lambda functions.

Answer: AExplanation: Edge-Optimized API: Designed for global users by routing requests through CloudFront's edge locations, reducing latency. Content Encoding: Enabling content encoding compresses data, further optimizing performance by decreasing payload size. Caching: Adding API Gateway caching reduces the number of calls to Lambda and database backends, improving latency. Reserved Concurrency: Although useful, this does not directly affect latency for transferring static and dynamic content.

QUESTION 1170A law firm needs to make hundreds of files readable for the general public. The law firm must prevent members of the public from modifying or deleting the files before a specified future date. Which solution will meet these requirements MOST securely?  
A. Upload the files to an Amazon S3 bucket that is configured for static website hosting. Grant read-only IAM permissions to any AWS principals that access the S3 bucket until the specified date.  
B. Create a new Amazon S3 bucket. Enable S3 Versioning. Use S3 Object Lock and set a retention period based on the specified date. Create an Amazon CloudFront distribution to serve content from the bucket. Use an S3 bucket policy to restrict access to the CloudFront origin access control (OAC).  
C. Create a new Amazon S3 bucket. Enable S3 Versioning. Configure an event trigger to run an AWS Lambda function if a user modifies or deletes an object. Configure the Lambda function to replace the modified or deleted objects with the original versions of the objects from a private S3 bucket.  
D. Upload the files to an Amazon S3 bucket that is configured for static website hosting. Select the folder that contains the files. Use S3 Object Lock with a retention period based on the specified date. Grant read-only IAM permissions to any AWS principals that access the S3 bucket.

Answer: BExplanation: S3 Object Lock: Enables Write Once Read Many (WORM) protection for data, preventing objects from being deleted or modified for a set retention period. S3 Versioning: Helps maintain object versions and ensures a recovery path for accidental overwrites.

CloudFront Distribution: Ensures secure and efficient public access by serving content through an edge-optimized delivery network while protecting S3 data with OAC. Bucket Policy for OAC: Restricts public access to only the CloudFront origin, ensuring maximum security.

QUESTION 1171A media company hosts a web application on AWS. The application gives users the ability to upload and view videos. The application stores the videos in an Amazon S3 bucket. The company wants to ensure that only authenticated users can upload videos. Authenticated users must have the ability to upload videos only within a specified time frame after authentication. Which solution will meet these requirements with the LEAST operational overhead?  
A. Configure the application to generate IAM temporary security credentials for authenticated users.  
B. Create an AWS Lambda function that generates pre-signed URLs when a user authenticates.  
C. Develop a custom authentication service that integrates with Amazon Cognito to control and log direct S3 bucket access through the application.  
D. Use AWS Security Token Service (AWS STS) to assume a pre-defined IAM role that grants authenticated users temporary permissions to upload videos directly to the S3 bucket.

Answer: BExplanation: Pre-Signed URLs: Allow temporary access to S3 buckets, making it easy to manage time-limited upload permissions without complex operational overhead. Lambda for Automation: Automatically generates and provides pre-signed URLs when users authenticate, minimizing manual steps and code complexity. Least Operational Overhead: Requires no custom authentication service or deep integration with STS or Cognito.

QUESTION 1172A company has a large fleet of vehicles that are equipped with internet connectivity to send telemetry to the company. The company receives over 1 million data points every 5 minutes from the vehicles. The company uses the data in machine learning (ML) applications to predict vehicle maintenance needs and to preorder parts. The

company produces visual reports based on the captured data. The company wants to migrate the telemetry ingestion, processing, and visualization workloads to AWS. Which solution will meet these requirements?A. Use Amazon Timestream for LiveAnalytics to store the data points. Grant Amazon SageMaker permission to access the data for processing. Use Amazon QuickSight to visualize the data.B. Use Amazon DynamoDB to store the data points. Use DynamoDB Connector to ingest data from DynamoDB into Amazon EMR for processing. Use Amazon QuickSight to visualize the data.C. Use Amazon Neptune to store the data points. Use Amazon Kinesis Data Streams to ingest data from Neptune into an AWS Lambda function for processing. Use Amazon QuickSight to visualize the data.D. Use Amazon Timestream to for LiveAnalytics to store the data points. Grant Amazon SageMaker permission to access the data for processing. Use Amazon Athena to visualize the data.

Answer: AExplanation:Amazon Timestream: Purpose-built time series database optimized for telemetry and IoT data ingestion and analytics. Amazon SageMaker: Provides ML capabilities for predictive maintenance workflows. Amazon QuickSight: Efficiently generates interactive, real-time visual reports from Timestream data. Optimized for Scale: Timestream efficiently handles large-scale telemetry data with time-series indexing and queries.

QUESTION 1173

A company runs an application on Amazon EC2 instances. The instances need to access an Amazon RDS database by using specific credentials. The company uses AWS Secrets Manager to contain the credentials the EC2 instances must use. Which solution will meet this requirement?A. Create an IAM role, and attach the role to each EC2 instance profile. Use an identity-based policy to grant the new IAM role access to the secret that contains the database credentials.B. Create an IAM user, and attach the user to each EC2 instance profile. Use a resource-based policy to grant the new IAM user access to the secret that contains the database credentials.C. Create a resource-based policy for the secret that contains the database credentials. Use EC2 Instance Connect to access the secret.D. Create an identity-based policy for the secret that contains the database credentials. Grant direct access to the EC2 instances.

Answer: AExplanation: IAM Role: Attaching an IAM role to an EC2 instance profile is a secure way to manage permissions without embedding credentials. AWS Secrets Manager: Grants controlled access to database credentials and automatically rotates secrets if configured. Identity-Based Policy: Ensures the IAM role only has access to specific secrets, enhancing security.

QUESTION 1174

A company stores petabytes of historical medical information on premises. The company has a process to manage encryption of the data to comply with regulations. The company needs a cloud-based solution for data backup, recovery, and archiving. The company must retain control over the encryption key material. Which combination of solutions will meet these requirements? (Choose two.)

A. Create an AWS Key Management Service (AWS KMS) key without key material. Import the company's key material into the KMS key.

B. Create an AWS Key Management Service (AWS KMS) encryption key that contains key material generated by AWS KMS.

C. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA) storage. Use S3 Bucket Keys with AWS Key Management Service (AWS KMS) keys.

D. Store the data in an Amazon S3 Glacier storage class. Use server-side encryption with customer-provided keys (SSE-C).

E. Store the data in AWS Snowball devices. Use server-side encryption with AWS KMS keys (SSE-KMS).

Answer: AExplanation: Option A: Importing customer-managed keys into AWS KMS ensures that encryption key material remains under the company's control. Option D: S3 Glacier with server-side encryption using customer-provided keys (SSE-C) complies with the need for controlled encryption and provides cost-effective storage for backups.

QUESTION 1175

A company is developing a social media application. The company anticipates rapid and unpredictable growth in users and data volume. The application needs to handle a continuous high volume of user requests. User requests include long-running processes that store large amounts of user-generated content and user profiles in a relational format. The processes must run in a specific order. The company requires an architecture that can scale resources to meet demand spikes without downtime or performance degradation. The company must ensure that the components of the application can evolve independently without affecting other parts of the system. Which combination of AWS services will meet these requirements?A. Deploy the application on Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Use Amazon RDS as the database. Use Amazon Simple Queue Service (Amazon SQS) to decouple message processing between components.

B. Deploy the application on Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Use Amazon RDS as the database. Use Amazon Simple Notification Service (Amazon SNS) to decouple message processing between components.

C. Use Amazon DynamoDB as the database. Use AWS Lambda functions to implement the application. Configure Amazon DynamoDB Streams to invoke the Lambda functions. Use AWS Step Functions to manage workflows between services.

D. Use an AWS Elastic Beanstalk environment with auto scaling to deploy the application. Use Amazon RDS as the database. Use Amazon Simple Notification Service (Amazon SNS) to decouple message processing between components.

Answer: AExplanation: ECS with Fargate: Allows containerized workloads to scale rapidly without managing underlying servers, handling unpredictable growth effectively. RDS for Relational Data: Manages large relational datasets efficiently while supporting high availability. SQS for Decoupling: Ensures message processing occurs in a specific order, decoupling application components and allowing independent evolution.

QUESTION 1176

A company plans to use AWS to run high-performance computing (HPC) workloads and analytics

workloads. The company will run HPC workloads on Amazon EC2 instances. The workloads require a high-performance file system that can scale to millions of input/output operations per second (IOPS). Which combination of steps will meet these requirements?

(Choose two.)A. Use Amazon Elastic File System (Amazon EFS) as a high-performance file system.B. Use Amazon FSx for Lustre as a high-performance file system.C. Create an Auto Scaling group of Amazon EC2 instances. Use Reserved Instances.

Configure a spread placement group. Use AWS Batch to run the analytics workloads.D. Use Mountpoint for Amazon S3 as a high-performance file system.E. Create an Auto Scaling group of Amazon EC2 instances. Use a mix of On-Demand Instances,

Reserved Instances, and Spot Instances. Configure a cluster placement group. Use Amazon EMR to run the analytics

workloads.

Answer: BEExplanation: Option B (Amazon FSx for Lustre): FSx for Lustre is optimized for high-performance file

systems required by HPC workloads, scaling to millions of IOPS and supporting parallelized data access. Option E (Cluster

Placement Group with Auto Scaling): A cluster placement group ensures low- latency communication between EC2 instances,

critical for HPC workloads. Amazon EMR simplifies running large-scale analytics jobs.

QUESTION 1177A company provides a trading platform to customers. The platform uses an Amazon API Gateway REST API, AWS Lambda functions, and an Amazon

DynamoDB table. Each trade that the platform processes invokes a Lambda function that stores the trade data in Amazon

DynamoDB. The company wants to ingest trade data into a data lake in Amazon S3 for near real-time analysis. Which solution will

meet these requirements with the LEAST operational overhead?

A. Use Amazon DynamoDB Streams to capture the trade data changes. Configure DynamoDB Streams to invoke a Lambda function that writes the data to Amazon S3.

B. Use Amazon DynamoDB Streams to capture the trade data changes. Configure DynamoDB Streams to invoke a Lambda function that writes the data to Amazon Data Firehose.

Write the data from Data Firehose to Amazon S3.

C. Enable Amazon Kinesis Data Streams on the DynamoDB table to capture the trade data changes. Configure Kinesis Data Streams to invoke a Lambda function that writes the data to Amazon S3.

D. Enable Amazon Kinesis Data Streams on the DynamoDB table to capture the trade data changes. Configure a data stream to be the input for Amazon Data Firehose. Write the data from Data Firehose to Amazon S3.

Answer: AExplanation: DynamoDB Streams: Captures real-time changes in DynamoDB tables and allows integration with Lambda for processing the

changes. Minimal Operational Overhead: Using a Lambda function directly to write data to S3 ensures simplicity and reduces the complexity of the pipeline.

QUESTION 1178A company has a large amount of data in an Amazon DynamoDB table. A large batch of data is appended to the table once each day. The company wants a solution that will make all the existing and future data in

DynamoDB available for analytics on a long-term basis. Which solution meets these requirements with the LEAST operational

overhead?

A. Configure DynamoDB incremental exports to Amazon S3.

B. Configure Amazon DynamoDB Streams to write records to Amazon S3.

C. Configure Amazon EMR to copy DynamoDB data to Amazon S3.

D. Configure Amazon EMR to copy DynamoDB data to Hadoop Distributed File System (HDFS).

Answer: AExplanation: Incremental Exports: Exporting DynamoDB data directly to Amazon S3 provides an automated, serverless way to make data available for analytics without operational overhead.

Analytics-Friendly Storage: Amazon S3 supports long-term analytics workloads and can integrate with tools like Athena or

QuickSight.

QUESTION 1179A company is building a serverless application to process large video files that users upload. The

application performs multiple tasks to process each video file. Processing can take up to 30 minutes for the largest files. The

company needs a scalable architecture to support the processing application. Which solution will meet these requirements?

A. Store the uploaded video files in Amazon Elastic File System (Amazon EFS). Configure a schedule in Amazon EventBridge Scheduler to

invoke an AWS Lambda function periodically to check for new files. Configure the Lambda function to perform all the processing

tasks.

B. Store the uploaded video files in Amazon Elastic File System (Amazon EFS). Configure an Amazon EFS event

notification to start an AWS Step Functions workflow that uses AWS Fargate tasks to perform the processing tasks.

C. Store the uploaded video files in Amazon S3. Configure an Amazon S3 event notification to send an event to Amazon EventBridge when a

user uploads a new video file. Configure an AWS Step Functions workflow as a target for an EventBridge rule. Use the workflow to

manage AWS Fargate tasks to perform the processing tasks.

D. Store the uploaded video files in Amazon S3. Configure an Amazon S3 event notification to invoke an AWS Lambda function when a user uploads a new video file. Configure the Lambda

function to perform all the processing tasks.

Answer: CExplanation: Scalability: The solution must scale as video files are uploaded.

Long-running tasks: Processing tasks can take up to 30 minutes. AWS Lambda has a maximum execution time of 15 minutes, which

rules out options that involve Lambda performing all the processing.

Serverless and event-driven architecture: Ensures cost-effectiveness and high availability.

QUESTION 1180A company is migrating a new application from an on-premises data center to a new VPC in the AWS Cloud. The company has multiple AWS accounts and VPCs that share many subnets and applications.

The company wants to have fine-grained access control for the new application. The company wants to ensure that all network

resources across accounts and VPCs that are granted permission to access the new application can access the application. Which

solution will meet these requirements?

A. Set up a VPC peering connection for each VPC that needs access to the new application

VPC.Update route tables in each VPC to enable connectivity.B. Deploy a transit gateway in the account that hosts the new application. Share the transit gateway with each account that needs to connect to the application. Update route tables in the VPC that hosts the new application and in the transit gateway to enable connectivity.C. Use an AWS PrivateLink endpoint service to make the new application accessible to other VPCs. Control access to the application by using an endpoint policy.D. Use an Application Load Balancer (ALB) to expose the new application to the internet.Configure authentication and authorization processes to ensure that only specified VPCs can access the application.Answer: CExplanation: AWS PrivateLink is the most suitable solution for providing fine-grained access control while allowing multiple VPCs, potentially across multiple accounts, to access the new application. This approach offers the following advantages:Fine-grained control: Endpoint policies can restrict access to specific services or principals. No need for route table updates: Unlike VPC peering or transit gateways, AWS PrivateLink does not require complex route table management.Scalable architecture: PrivateLink scales to support traffic from multiple VPCs. Secure connectivity: Ensures private connectivity over the AWS network, without exposing resources to the internet.QUESTION 1181A solutions architect needs to implement a solution that can handle up to 5,000 messages per second. The solution must publish messages as events to multiple consumers. The messages are up to 500 KB in size. The message consumers need to have the ability to use multiple programming languages to consume the messages with minimal latency. The solution must retain published messages for more than 3 months. The solution must enforce strict ordering of the messages. Which solution will meet these requirements?A. Publish messages to an Amazon Kinesis Data Streams data stream. Enable enhanced fan-out. Ensure that consumers ingest the data stream by using dedicated throughput.B. Publish messages to an Amazon Simple Notification Service (Amazon SNS) topic. Ensure that consumers use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to subscribe to the topic.C. Publish messages to Amazon EventBridge. Allow each consumer to create rules to deliver messages to the consumer's own target.D. Publish messages to an Amazon Simple Notification Service (Amazon SNS) topic. Ensure that consumers use Amazon Data Firehose to subscribe to the topic.Answer: AExplanation: Amazon Kinesis Data Streams is the best choice for this scenario:Message throughput: Kinesis Data Streams supports high throughput with enhanced fan-out and dedicated throughput for consumers.Large message size: Supports message sizes up to 1 MB, meeting the 500 KB requirement. Message retention: Data streams can retain messages for up to 365 days. Strict ordering: Guarantees message ordering within shards.QUESTION 1182A healthcare provider is planning to store patient data on AWS as PDF files. To comply with regulations, the company must encrypt the data and store the files in multiple locations. The data must be available for immediate access from any environment. Which solution will meet these requirements?A. Store the files in an Amazon S3 bucket. Use the Standard storage class. Enable server-side encryption with Amazon S3 managed keys (SSE-S3) on the bucket. Configure cross-Region replication on the bucket.B.

Store the files in an Amazon Elastic File System (Amazon EFS) volume. Use an AWS KMS managed key to encrypt the EFS volume. Use AWS DataSync to replicate the EFS volume to a second AWS Region.C. Store the files in an Amazon Elastic Block Store (Amazon EBS) volume. Configure AWS Backup to back up the volume on a regular schedule. Use an AWS KMS key to encrypt the backups.D. Store the files in an Amazon S3 bucket. Use the S3 Glacier Flexible Retrieval storage class. Ensure that all PDF files are encrypted by using client-side encryption before the files are uploaded. Configure cross-Region replication on the bucket.Answer: AExplanation: Amazon S3 with the Standard storage class is the best solution:Encryption: SSE-S3 ensures server-side encryption of the data, meeting compliance requirements. Immediate access: The Standard storage class provides low-latency and high-throughput access to data. Multi-location storage: Cross-Region replication ensures data is stored in multiple AWS Regions for redundancy.QUESTION 1183A company wants to use an API to translate text from one language to another. The API must receive an HTTP header value and pass the value to an embedded library. The API translates documents in 6 minutes. The API requires a custom authorization mechanism. Which solution will meet these requirements?A. Configure an Amazon API Gateway REST API with AWS\_PROXY integration to synchronously call an AWS Lambda function to perform translations.B. Configure an AWS Lambda function with a Lambda function URL to synchronously call a second function to perform translations.C. Configure an Amazon API Gateway REST API with AWS\_PROXY integration to asynchronously call an AWS Lambda function to perform translations.D. Configure an Amazon API Gateway REST API with HTTP\_PROXY integration to synchronously call a web endpoint that is hosted on an EC2 instance.Answer: AExplanation: The AWS\_PROXY integration with Amazon API Gateway allows the API to invoke a Lambda function synchronously, making it a suitable solution for the custom authorization mechanism and text translation use case.Synchronous Invocation: The API Gateway REST API with AWS\_PROXY integration enables synchronous processing of HTTP requests and responses, which is required for document translation. Custom Authorization: API Gateway supports custom authorizers for fine-grained access control. Lambda Function Execution: Although Lambda's execution time limit is 15 minutes, this is sufficient for the 6-minute document translation requirement.QUESTION 1184A company uses Amazon S3 to store customer data that contains personally identifiable information (PII) attributes. The company needs to make the customer

information available to company resources through an AWS Glue Catalog. The company needs to have fine-grained access control for the data so that only specific IAM roles can access the PII data. Which solution will meet these requirements? A. Create one IAM policy that grants access to PII. Create a second IAM policy that grants access to non-PII data. Assign the PII policy to the specified IAM roles. B. Create one IAM role that grants access to PII. Create a second IAM role that grants access to non-PII data. Assign the PII policy to the specified IAM roles. C. Use AWS Lake Formation to provide the specified IAM roles access to the PII data. D. Use AWS Glue to create one view for PII data. Create a second view for non-PII data. Provide the specified IAM roles access to the PII view. Answer: C Explanation: AWS Lake Formation is designed for managing fine-grained access control to data in an efficient manner. Granular Permissions: Lake Formation allows column-level, row-level, and table-level access controls, which can precisely define access to PII data. Integration with AWS Glue Catalog: Lake Formation natively integrates with AWS Glue for seamless data cataloging and access control. Operational Efficiency: Centralized access control policies minimize the need for separate IAM roles or policies.

**QUESTION 1185** A company stores 5 PB of archived data on physical tapes. The company needs to preserve the data for another 10 years. The data center that stores the tapes has a 10 Gbps Direct Connect connection to an AWS Region. The company wants to migrate the data to AWS within the next 6 months. Which solution will meet these requirements? A. Read the data from the tapes on premises. Use local storage to stage the data. Use AWS DataSync to migrate the data to Amazon S3 Glacier Flexible Retrieval storage. B. Use an on-premises backup application to read the data from the tapes. Use the backup application to write directly to Amazon S3 Glacier Deep Archive storage. C. Order multiple AWS Snowball Edge devices. Copy the physical tapes to virtual tapes on the Snowball Edge devices. Ship the Snowball Edge devices to AWS. Create an S3 Lifecycle policy to move the tapes to Amazon S3 Glacier Instant Retrieval storage. D. Configure an on-premises AWS Storage Gateway Tape Gateway. Create virtual tapes in the AWS Cloud. Use backup software to copy the physical tapes to the virtual tapes. Move the virtual tapes to Amazon S3 Glacier Deep Archive storage. Answer: D Explanation: The company's requirements are to migrate 5 PB of data from physical tapes to AWS within 6 months, preserve the data for 10 years, and ensure cost efficiency. AWS Storage Gateway Tape Gateway is purpose-built for such use cases, as it seamlessly integrates with backup applications and provides virtual tape storage in Amazon S3 Glacier Deep Archive.

**QUESTION 1186** A company is using microservices to build an ecommerce application on AWS. The company wants to preserve customer transaction information after customers submit orders. The company wants to store transaction data in an Amazon Aurora database. The company expects sales volumes to vary throughout each year. Which solution will meet these requirements? A. Use an Amazon API Gateway REST API to invoke an AWS Lambda function to send transaction data to the Aurora database. Send transaction data to an Amazon Simple Queue Service (Amazon SQS) queue that has a dead-letter queue. Use a second Lambda function to read from the SQS queue and to update the Aurora database. B. Use an Amazon API Gateway HTTP API to send transaction data to an Application Load Balancer (ALB). Use the ALB to send the transaction data to Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use ECS tasks to store the data in Aurora database. C. Use an Application Load Balancer (ALB) to route transaction data to Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon EKS to send the data to the Aurora database. D. Use Amazon Data Firehose to send transaction data to Amazon S3. Use AWS Database Migration Service (AWS DMS) to migrate the data from Amazon S3 to the Aurora database. Answer: A Explanation: The solution must handle variable sales volumes, preserve transaction information, and store data in an Amazon Aurora database with minimal operational overhead. Using API Gateway, AWS Lambda, and SQS is the best option because it provides scalability, reliability, and resilience.

**QUESTION 1187** A company has an application that receives and processes purchase orders. The application supports only XML data. The company needs to configure the application to accept orders in JSON format. The company does not want to modify the application. A solutions architect is using an Amazon API Gateway HTTP API to create a new purchase order API. The solutions architect needs to modify the application DNS record to point to the new HTTP API. Which solution will meet these requirements? A. Use an HTTP proxy integration to pass XML requests to the application. For JSON requests, use API Gateway mappings to convert the purchase orders to XML. Use an AWS Lambda function that is integrated with API Gateway to call the application. B. Use an HTTP proxy integration to pass XML requests to the application. For JSON requests, use an AWS Lambda function that is integrated with API Gateway to convert the purchase orders from JSON to XML and to call the application. C. Use an HTTP custom integration to pass XML requests to the application. For JSON requests, use API Gateway mappings to convert the purchase orders to XML. Use an AWS Lambda function that is integrated with API Gateway to call the application. D. Use an HTTP custom integration to pass XML requests to the application. For JSON requests, use an AWS Lambda function that is integrated with API Gateway to convert the purchase orders to JSON and to call the application. Answer: B Explanation: HTTP Proxy Integration: Passes XML requests directly to the application, which already supports XML. JSON Conversion: An AWS Lambda function converts JSON requests to XML and calls the application. API Gateway: Acts as a front end to handle JSON requests and integrates seamlessly with Lambda for the transformation process.

**QUESTION 1188** A company stores

data for multiple business units in a single Amazon S3 bucket that is in the company's payer AWS account. To maintain data isolation, the business units store data in separate prefixes in the S3 bucket by using an S3 bucket policy. The company plans to add a large number of dynamic prefixes. The company does not want to rely on a single S3 bucket policy to manage data access at scale. The company wants to develop a secure access management solution in addition to the bucket policy to enforce prefix-level data isolation. Which solution will meet these requirements?

- A. Configure the S3 bucket policy to deny s3:GetObject permissions for all users. Configure the bucket policy to allow s3:\* access to individual business units.
- B. Enable default encryption on the S3 bucket by using server-side encryption with Amazon S3 managed keys (SSE-S3).
- C. Configure resource-based permissions on the S3 bucket by creating an S3 access point for each business unit.
- D. Use pre-signed URLs to provide access to the S3 bucket.

Answer: C  
Explanation: S3 Access Points: Provide scalable management of access to large datasets with specific permissions for individual prefixes. Dynamic Prefixes: Access points simplify managing access to a growing number of prefixes without relying solely on a single bucket policy. Fine-Grained Control: Resource-based permissions on access points enforce prefix-level isolation effectively.

QUESTION 1189 A company is migrating an online marketplace application from a mainframe system to an Auto Scaling group of Amazon EC2 instances. The EC2 instances access an Amazon Aurora cluster. The application requires a scalable, persistent caching solution to store the results of in-progress transactions and SQL queries. Which solution will meet these requirements?

- A. Use an Amazon ElastiCache (Redis OSS) cluster to serve transaction and query results.
- B. Use an Amazon CloudFront distribution with an Amazon S3 bucket as the origin to cache the transactions. Add an Amazon EC2 instance store volume to the EC2 instances for query result caching.
- C. Use an Amazon ElastiCache (Memcached) cluster to serve transaction and query results.
- D. Use an Amazon ElastiCache (Redis OSS) cluster to cache the transactions. Add an Amazon EC2 instance store volume to the EC2 instances for query result caching.

Answer: A  
Explanation: ElastiCache for Redis: Provides persistent, scalable caching for in-progress transactions and SQL queries. Redis supports data durability and advanced features, making it suitable for transactional workloads. Integration with Aurora: Easily integrates with the Aurora cluster to improve query performance.

QUESTION 1190 A company is launching a new application that will be hosted on Amazon EC2 instances. A solutions architect needs to design a solution that does not allow public IPv4 access that originates from the internet. However, the solution must allow the EC2 instances to make outbound IPv4 internet requests. Which solution will meet these requirements?

- A. Deploy a NAT gateway in public subnets in both Availability Zones. Create and configure one route table for each private subnet.
- B. Deploy an internet gateway in public subnets in both Availability Zones. Create and configure a shared route table for the private subnets.
- C. Deploy a NAT gateway in public subnets in both Availability Zones. Create and configure a shared route table for the private subnets.
- D. Deploy an egress-only internet gateway in public subnets in both Availability Zones. Create and configure one route table for each private subnet.

Answer: C  
Explanation: NAT Gateway: Allows private subnets to access the internet for outbound requests while preventing inbound connections. High Availability: Deploying NAT gateways in both AZs ensures fault tolerance.

Shared Route Table: Simplifies routing configuration for private subnets.

QUESTION 1191 A company hosts an application on Amazon EC2 instances that are part of a target group behind an Application Load Balancer (ALB). The company has attached a security group to the ALB. During a recent review of application logs, the company found many unauthorized login attempts from IP addresses that belong to countries outside the company's normal user base. The company wants to allow traffic only from the United States and Australia. Which solution will meet these requirements?

- A. Edit the default network ACL to block IP addresses from outside of the allowed countries.
- B. Create a geographic match rule in AWS WAF. Attach the rule to the ALB.
- C. Configure the ALB security group to allow the IP addresses of company employees. Edit the default network ACL to block IP addresses from outside of the allowed countries.
- D. Use a host-based firewall on the EC2 instances to block IP addresses from outside of the allowed countries. Configure the ALB security group to allow the IP addresses of company employees.

Answer: B  
Explanation: AWS WAF: Provides a simple way to create geographic match rules to block or allow traffic based on country IP ranges. Least Operational Overhead: Attaching the WAF rule to the ALB ensures centralized control without modifying ACLs or instance firewalls.

QUESTION 1192 A company has an ecommerce application that users access through multiple mobile apps and web applications. The company needs a solution that will receive requests from the mobile apps and web applications through an API. Request traffic volume varies significantly throughout each day. Traffic spikes during sales events. The solution must be loosely coupled and ensure that no requests are lost. Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB). Create an AWS Elastic Beanstalk endpoint to process the requests. Add the Elastic Beanstalk endpoint to the target group of the ALB.
- B. Set up an Amazon API Gateway REST API with an integration to an Amazon Simple Queue Service (Amazon SQS) queue. Configure a dead-letter queue. Create an AWS Lambda function to poll the queue to process the requests.
- C. Create an Application Load Balancer (ALB). Create an AWS Lambda function to process the requests. Add the Lambda function as a target of the ALB.
- D. Set up an Amazon API Gateway HTTP API with an integration to an Amazon Simple Notification Service (Amazon SNS) topic.

SNS) topic. Create an AWS Lambda function to process the requests. Subscribe the function to the SNS topic to process the requests.

Answer: BExplanation:Amazon SQS: Ensures no requests are lost, even during traffic spikes. API Gateway: Handles dynamic traffic patterns efficiently, integrating with SQS for asynchronous processing. Lambda: Polls the queue and processes requests in a serverless and scalable manner. Dead-Letter Queue (DLQ): Ensures failed messages are retried or logged for debugging.

QUESTION 1193A company is developing a new application that will run on Amazon EC2 instances. The application needs to access multiple AWS services. The company needs to ensure that the application will not use long-term access keys to access AWS services. Which solution will meet these requirements?

A. Create an IAM user. Assign the IAM user to the application. Create programmatic access keys for the IAM user. Embed the access keys in the application code.

B. Create an IAM user that has programmatic access keys. Store the access keys in AWS Secrets Manager. Configure the application to retrieve the keys from Secrets Manager when the application runs.

C. Create an IAM role that can access AWS Systems Manager Parameter Store. Associate the role with each EC2 instance profile. Create IAM access keys for the AWS services, and store the keys in Parameter Store. Configure the application to retrieve the keys from Parameter Store when the application runs.

D. Create an IAM role that has permissions to access the required AWS services. Associate the IAM role with each EC2 instance profile.

Answer: DExplanation: IAM Roles with Instance Profiles: Allow applications to access AWS services securely without hardcoding long-term access keys. Short-Term Credentials: IAM roles issue short-term credentials dynamically managed by AWS.

QUESTION 1194A company is developing a containerized web application that needs to be highly available and scalable. The application requires access to GPU resources. Which solution will meet these requirements?

A. Package the application as an AWS Lambda function in a container image. Use Lambda to run the containerized application on a runtime with GPU access.

B. Deploy the application container to Amazon Elastic Kubernetes Service (Amazon EKS). Use AWS Fargate to manage compute resources and access to GPU resources.

C. Deploy the application container to Amazon Elastic Container Registry (Amazon ECR). Use Amazon ECR to run the containerized application with an attached GPU.

D. Run the application on Amazon EC2 instances from a GPU instance family by using Amazon Elastic Container Service (Amazon ECS) for orchestration.

Answer: DExplanation: GPU Access: Only EC2 instances in the GPU family (e.g., P2, P3) can provide GPU resources. ECS Orchestration: Simplifies container deployment and management.

QUESTION 1195A solutions architect needs to secure an Amazon API Gateway REST API. Users need to be able to log in to the API by using common external social identity providers (IdPs). The social IdPs must use standard authentication protocols such as SAML or OpenID Connect (OIDC). The solutions architect needs to protect the API against attempts to exploit application vulnerabilities. Which combination of steps will meet these security requirements? (Choose two.)

A. Create an AWS WAF web ACL that is associated with the REST API. Add the appropriate managed rules to the ACL.

B. Subscribe to AWS Shield Advanced. Enable DDoS protection. Associate Shield Advanced with the REST API.

C. Create an Amazon Cognito user pool with a federation to the social IdPs. Integrate the user pool with the REST API.

D. Create an API key in API Gateway. Associate the API key with the REST API.

E. Create an IP address filter in AWS WAF that allows only the social IdPs. Associate the filter with the web ACL and the API.

Answer: ACQUESTION 1196A finance company uses backup software to back up its data to physical tape storage on-premises. To comply with regulations, the company needs to store the data for 7 years. The company must be able to restore archived data within one week when necessary. The company wants to migrate the backup data to AWS to reduce costs. The company does not want to change the current backup software. Which solution will meet these requirements MOST cost-effectively?

A. Use AWS Storage Gateway Tape Gateway to copy the data to virtual tapes. Use AWS DataSync to migrate the virtual tapes to the Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Change the target of the backup software to S3 Standard-IA.

B. Convert the physical tapes to virtual tapes. Use AWS DataSync to migrate the virtual tapes to Amazon S3 Glacier Flexible Retrieval. Change the target of the backup software to the S3 Glacier Flexible Retrieval.

C. Use AWS Storage Gateway Tape Gateway to copy the data to virtual tapes. Migrate the virtual tapes to Amazon S3 Glacier Deep Archive. Change the target of the backup software to the virtual tapes.

D. Convert the physical tapes to virtual tapes. Use AWS Snowball Edge storage-optimized devices to migrate the virtual tapes to Amazon S3 Glacier Flexible Retrieval. Change the target of the backup software to S3 Glacier Flexible Retrieval.

Answer: CExplanation: AWS Storage Gateway Tape Gateway provides a seamless way to migrate backup data to AWS without requiring changes to the backup software. Migrating to S3 Glacier Deep Archive ensures long-term, cost-effective storage for data that rarely needs retrieval. Option A:S3 Standard-IA is more expensive than Glacier for long-term storage. Option B and D:Glacier Flexible Retrieval is costlier than Glacier Deep Archive for archival use cases with low retrieval frequency.

QUESTION 1197A company is designing a new application that uploads files to an Amazon S3 bucket. The uploaded files are processed to extract metadata. Processing must take less than 5 seconds. The volume and frequency of the uploads vary from a few files each hour to hundreds of concurrent uploads. Which solution will meet these requirements MOST cost-effectively?

A. Configure AWS CloudTrail trails to log Amazon S3 API calls. Use AWS AppSync to process the files.

B. Configure a new object

created S3 event notification within the bucket to invoke an AWS Lambda function to process the files.C. Configure Amazon Kinesis Data Streams to deliver the files to the S3 bucket. Invoke an AWS Lambda function to process the files.D. Deploy an Amazon EC2 instance. Create a script that lists all files in the S3 bucket and processes new files. Use a cron job that runs every minute to run the script.

**Answer:** BExplanation: Using S3 event notifications to trigger AWS Lambda for file processing is a cost-effective and serverless solution. Lambda scales automatically with upload volume, and processing each file takes less than 5 seconds, fitting within Lambda's execution time.

**QUESTION 1198** A solutions architect is designing the network architecture for an application that runs on Amazon EC2 instances in an Auto Scaling group. The application needs to access data that is in Amazon S3 buckets. Traffic to the S3 buckets must not use public IP addresses. The solutions architect will deploy the application in a VPC that has public and private subnets. Which solutions will meet these requirements? (Choose two.)

A. Deploy the EC2 instances in a private subnet. Configure a default route to an egress-only internet gateway.

B. Deploy the EC2 instances in a public subnet. Create a gateway endpoint for Amazon S3. Associate the endpoint with the subnet's route table.

C. Deploy the EC2 instances in a public subnet. Create an interface endpoint for Amazon S3. Configure DNS hostnames and DNS resolution for the VPC.

D. Deploy the EC2 instances in a private subnet. Configure a default route to a NAT gateway in a public subnet.

E. Deploy the EC2 instances in a private subnet. Configure a default route to a customer gateway.

**Answer:** BD

**QUESTION 1199** A company is building a serverless application to process orders from an ecommerce site. The application needs to handle bursts of traffic during peak usage hours and to maintain high availability. The orders must be processed asynchronously in the order the application receives them. Which solution will meet these requirements?

A. Use an Amazon Simple Notification Service (Amazon SNS) topic to receive orders. Use an AWS Lambda function to process the orders.

B. Use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to receive orders. Use an AWS Lambda function to process the orders.

C. Use an Amazon Simple Queue Service (Amazon SQS) standard queue to receive orders. Use AWS Batch jobs to process the orders.

D. Use an Amazon Simple Notification Service (Amazon SNS) topic to receive orders. Use AWS Batch jobs to process the orders.

**Answer:** BExplanation: Amazon SQS FIFO queues ensure that orders are processed in the exact order received and maintain message deduplication. AWS Lambda scales automatically, handling bursts and maintaining high availability in a cost-effective manner.

**QUESTION 1200** A company's reporting system delivers hundreds of .csv files to an Amazon S3 bucket each day. The company must convert these files to Apache Parquet format and must store the files in a transformed data bucket. Which solution will meet these requirements with the LEAST development effort?

A. Create an Amazon EMR cluster with Apache Spark installed. Write a Spark application to transform the data. Use EMR File System (EMRFS) to write files to the transformed data bucket.

B. Create an AWS Glue crawler to discover the data. Create an AWS Glue extract, transform, and load (ETL) job to transform the data. Specify the transformed data bucket in the output step.

C. Use AWS Batch to create a job definition with Bash syntax to transform the data and output the data to the transformed data bucket. Use the job definition to submit a job. Specify an array job as the job type.

D. Create an AWS Lambda function to transform the data and output the data to the transformed data bucket. Configure an event notification for the S3 bucket. Specify the Lambda function as the destination for the event notification.

**Answer:** BExplanation: AWS Glue provides a serverless ETL solution requiring minimal development. Glue supports conversion to Parquet with managed jobs and integrates with S3 for output.

**QUESTION 1201** A company is designing a new Amazon Elastic Kubernetes Service (Amazon EKS) deployment to host multi-tenant applications that use a single cluster. The company wants to ensure that each pod has its own hosted environment. The environments must not share CPU, memory, storage, or elastic network interfaces. Which solution will meet these requirements?

A. Use Amazon EC2 instances to host self-managed Kubernetes clusters. Use taints and tolerations to enforce isolation boundaries.

B. Use Amazon EKS with AWS Fargate. Use Fargate to manage resources and to enforce isolation boundaries.

C. Use Amazon EKS and self-managed node groups. Use taints and tolerations to enforce isolation boundaries.

D. Use Amazon EKS and managed node groups. Use taints and tolerations to enforce isolation boundaries.

**Answer:** BExplanation: AWS Fargate provides per-pod isolation for CPU, memory, storage, and networking, making it ideal for multi-tenant use cases.

**QUESTION 1202** A company is developing a public web application that needs to access multiple AWS services. The application will have hundreds of users who must log in to the application first before using the services. The company needs to implement a secure and scalable method to grant the web application temporary access to the AWS resources. Which solution will meet these requirements?

A. Create an IAM role for each AWS service that the application needs to access. Assign the roles directly to the instances that the web application runs on.

B. Create an IAM role that has the access permissions the web application requires. Configure the web application to use AWS Security Token Service (AWS STS) to assume the IAM role. Use STS tokens to access the required AWS services.

C. Use AWS IAM Identity Center to create a user pool that includes the application users. Assign access credentials to the web application users. Use the credentials to access the required AWS services.

D. Create an IAM user that has programmatic access keys for the AWS services. Store the access keys in AWS Systems Manager Parameter Store. Retrieve the access keys from Parameter Store. Use the

keys in the web application. Answer: BExplanation: AWS Security Token Service (STS) allows the web application to request temporary security credentials that grant access to AWS resources. These temporary credentials are secure and short-lived, reducing the risk of misuse. Using STS and IAM roles ensures scalability by enabling the application to dynamically assume roles with the required permissions for each AWS service. QUESTION 1203 A company wants to run big data workloads on Amazon EMR. The workloads need to process terabytes of data in memory. A solutions architect needs to identify the appropriate EMR cluster instance configuration for the workloads. Which solution will meet these requirements? A. Use a storage optimized instance for the primary node. Use compute optimized instances for core nodes and task nodes. B. Use a memory optimized instance for the primary node. Use storage optimized instances for core nodes and task nodes. C. Use a general purpose instance for the primary node. Use memory optimized instances for core nodes and task nodes. D. Use general purpose instances for the primary, core, and task nodes. Answer: CExplanation: Big data workloads that need to process terabytes of data in memory require memory-optimized instances for the core and task nodes to ensure sufficient memory for processing data efficiently. Primary Node: A general purpose instance is suitable because it manages cluster operations, including coordination and monitoring, and does not process data directly. Core and Task Nodes: These nodes handle data storage and processing. Memory-optimized instances are ideal because they provide high memory-to-CPU ratios, which is critical for in-memory big data workloads. QUESTION 1204 A solutions architect needs to optimize a large data analytics job that runs on an Amazon EMR cluster. The job takes 13 hours to finish. The cluster has multiple core nodes and worker nodes deployed on large, compute-optimized instances. After reviewing EMR logs, the solutions architect discovers that several nodes are idle for more than 5 hours while the job is running. The solutions architect needs to optimize cluster performance. Which solution will meet this requirement MOST cost-effectively? A. Increase the number of core nodes to ensure there is enough processing power to handle the analytics job without any idle time. B. Use the EMR managed scaling feature to automatically resize the cluster based on workload. C. Migrate the analytics job to a set of AWS Lambda functions. Configure reserved concurrency for the functions. D. Migrate the analytics job core nodes to a memory-optimized instance type to reduce the total job runtime. Answer: BExplanation: EMR managed scaling dynamically resizes the cluster by adding or removing nodes based on the workload. This feature helps minimize idle time and reduces costs by scaling the cluster to meet processing demands efficiently. QUESTION 1205 A solutions architect is building an Amazon S3 data lake for a company. The company uses Amazon Kinesis Data Firehose to ingest customer personally identifiable information (PII) and transactional data in near real-time to an S3 bucket. The company needs to mask all PII data before storing the data in the data lake. Which solution will meet these requirements? A. Create an AWS Lambda function to detect and mask PII. Invoke the function from Kinesis Data Firehose. B. Use Amazon Macie to scan the S3 bucket. Configure Macie to detect and mask PII. C. Enable server-side encryption (SSE) on the S3 bucket. D. Create an AWS Lambda function that integrates with AWS CloudHSM. Configure the function to detect and mask PII. Answer: AExplanation: Using a Lambda function as part of the Kinesis Data Firehose pipeline allows for real-time detection and masking of PII before data is written to S3. This ensures that PII is never stored in its raw form in the data lake. QUESTION 1206 An ecommerce company runs an application that uses an Amazon DynamoDB table in a single AWS Region. The company wants to deploy the application to a second Region. The company needs to support multi-active replication with low latency reads and writes to the existing DynamoDB table in both Regions. Which solution will meet these requirements in the MOST operationally efficient way? A. Create a DynamoDB global secondary index (GSI) for the existing table. Create a new table in the second Region. Convert the existing DynamoDB table to a global table. Specify the new table as the secondary table. B. Enable Amazon DynamoDB Streams for the existing table. Create a new table in the second Region. Create a new application that uses the DynamoDB Streams Kinesis Adapter and the Amazon Kinesis Client Library (KCL). Configure the new application to read data from the DynamoDB table in the first Region and to write the data to the new table in the second Region. C. Convert the existing DynamoDB table to a global table. Choose the appropriate second Region to achieve active-active write capabilities in both Regions. D. Enable Amazon DynamoDB Streams for the existing table. Create a new table in the second Region. Create an AWS Lambda function in the first Region that reads data from the table in the first Region and writes the data to the new table in the second Region. Set a DynamoDB stream as the input trigger for the Lambda function. Answer: CExplanation: Converting the existing DynamoDB table to a global table provides active-active replication and low-latency reads and writes in both Regions. DynamoDB global tables are specifically designed for multi-Region and multi-active use cases. QUESTION 1207 A company runs a critical public application on Amazon Elastic Kubernetes Service (Amazon EKS) clusters. The application has a microservices architecture. The company needs to implement a solution that collects, aggregates, and summarizes metrics and logs from the application in a centralized location. Which solution will meet these requirements in the MOST operationally efficient way? A. Run the Amazon CloudWatch agent in the existing EKS cluster. Use a CloudWatch dashboard to view the metrics and logs. B. Configure a data stream in Amazon Kinesis Data Streams. Use Amazon Kinesis Data Firehose to read events and to deliver the events to an Amazon

S3 bucket. Use Amazon Athena to view the events.C. Configure AWS CloudTrail to capture data events. Use Amazon OpenSearch Service to query CloudTrail.D. Configure Amazon CloudWatch Container Insights in the existing EKS cluster. Use a CloudWatch dashboard to view the metrics and logs.

Answer: DExplanation:Amazon CloudWatch Container Insights is designed for monitoring containerized environments like EKS. It provides native support for collecting and visualizing metrics and logs in a centralized location through CloudWatch dashboards, offering the most operationally efficient solution.

QUESTION 1208A company hosts an application that processes highly sensitive customer transactions on AWS. The application uses Amazon RDS as its database. The company manages its own encryption keys to secure the data in Amazon RDS. The company needs to update the customer-managed encryption keys at least once each year. Which solution will meet these requirements with the LEAST operational overhead?

A. Set up automatic key rotation in AWS Key Management Service (AWS KMS) for the encryption keys.

B. Configure AWS Key Management Service (AWS KMS) to alert the company to rotate the encryption keys annually.

C. Schedule an AWS Lambda function to rotate the encryption keys annually.

D. Create an AWS CloudFormation stack to run an AWS Lambda function that deploys new encryption keys once each year.

Answer: AExplanation:AWS KMS automatic key rotation is the simplest and most operationally efficient solution. Enabling automatic key rotation ensures that KMS automatically generates new key material for the key every year without requiring manual intervention.

QUESTION 1209A company is designing a web application with an internet-facing Application Load Balancer (ALB). The company needs the ALB to receive HTTPS web traffic from the public internet. The ALB must send only HTTPS traffic to the web application servers hosted on the Amazon EC2 instances on port 443. The ALB must perform a health check of the web application servers over HTTPS on port 8443. Which combination of configurations of the security group that is associated with the ALB will meet these requirements? (Choose three.)

A. Allow HTTPS inbound traffic from 0.0.0.0/0 for port 443.

B. Allow all outbound traffic to 0.0.0.0/0 for port 443.

C. Allow HTTPS outbound traffic to the web application instances for port 443.

D. Allow HTTPS inbound traffic from the web application instances for port 443.

E. Allow HTTPS outbound traffic to the web application instances for the health check on port 8443.

F. Allow HTTPS inbound traffic from the web application instances for the health check on port 8443.

Answer: ACEExplanation: Option A: The ALB must accept HTTPS traffic from the public internet. Allowing inbound traffic on port 443 from 0.0.0.0/0 enables this functionality.

Option C: The ALB must forward HTTPS traffic to the web application servers on port 443. Outbound traffic for port 443 must be allowed for this communication.

Option E: The ALB must perform health checks on the web application servers over HTTPS on port 8443. Outbound traffic for port 8443 must be allowed for this purpose.

QUESTION 1210A finance company has a web application that generates credit reports for customers. The company hosts the frontend of the web application on a fleet of Amazon EC2 instances that is associated with an Application Load Balancer (ALB). The application generates reports by running queries on an Amazon RDS for SQL Server database.

The company recently discovered that malicious traffic from around the world is abusing the application by submitting unnecessary requests. The malicious traffic is consuming significant compute resources.

The company needs to address the malicious traffic. Which solution will meet this requirement?

A. Use AWS WAF to create a web ACL. Associate the web ACL with the ALB. Update the web ACL to block IP addresses that are associated with malicious traffic.

B. Use AWS WAF to create a web ACL. Associate the web ACL with the ALB. Use the AWS WAF Bot Control managed rule feature.

C. Set up AWS Shield to protect the ALB and the database.

D. Use AWS WAF to create a web ACL. Associate the web ACL with the ALB. Configure the AWS WAF IP reputation rule.

Answer: BExplanation: The AWS WAF Bot Control managed rule is designed to automatically detect and mitigate bot traffic. This feature is particularly useful for addressing malicious traffic and conserving compute resources by filtering unnecessary requests at the ALB level.

QUESTION 1211A media company is launching a new product platform that artists from around the world can use to upload videos and images directly to an Amazon S3 bucket.

The company owns and maintains the S3 bucket. The artists must be able to upload files from personal devices without the need for AWS credentials or an AWS account. Which solution will meet these requirements MOST securely?

A. Enable cross-origin resource sharing (CORS) on the S3 bucket.

B. Turn off block public access for the S3 bucket.

Share the bucket URL to the artists to enable uploads without credentials.

C. Use an IAM role that has upload permissions for the S3 bucket to generate presigned URLs for S3 prefixes that are specific to each artist.

Share the URLs to the artists.

D. Create a web interface that uses an IAM role that has permission to upload and view objects in the S3 bucket.

Share the web interface URL to the artists.

Answer: CExplanation: Presigned URLs allow temporary, limited access to upload files to specific S3 prefixes. This ensures that artists can upload files securely without needing AWS credentials or accounts. Each artist receives a unique URL with permissions tied to the intended S3 location, and the URL can be configured to expire after a certain time, minimizing security risks.

QUESTION 1212A company is building a serverless application to process clickstream data from its website.

The clickstream data is sent to an Amazon Kinesis Data Streams data stream from the application web servers.

The company wants to enrich the clickstream data by joining the clickstream data with customer profile data from an Amazon Aurora Multi-AZ database.

The

company wants to use Amazon Redshift to analyze the enriched data. The solution must be highly available. Which solution will meet these requirements?

A. Use an AWS Lambda function to process and enrich the clickstream data. Use the same Lambda function to write the clickstream data to Amazon S3. Use Amazon Redshift Spectrum to query the enriched data in Amazon S3.

B. Use an Amazon EC2 Spot Instance to poll the data stream and enrich the clickstream data. Configure the EC2 instance to use the COPY command to send the enriched results to Amazon Redshift.

C. Use an Amazon Elastic Container Service (Amazon ECS) task with AWS Fargate Spot capacity to poll the data stream and enrich the clickstream data. Configure an Amazon EC2 instance to use the COPY command to send the enriched results to Amazon Redshift.

D. Use Amazon Kinesis Data Firehose to load the clickstream data from Kinesis Data Streams to Amazon S3. Use AWS Glue crawlers to infer the schema and populate the AWS Glue Data Catalog. Use Amazon Athena to query the raw data in Amazon S3.

Answer: A

Explanation: Option A is the best solution as it leverages AWS Lambda for serverless, scalable, and highly available processing and enrichment of clickstream data. Lambda can process the data in real-time, join it with the Aurora database data, and write the enriched results to Amazon S3. From S3, Amazon Redshift Spectrum can directly query the enriched data without needing to load the data into Redshift, enabling cost efficiency and high availability.

QUESTION 1213

A company wants to create an API to authorize users by using JSON Web Tokens (JWTs). The company needs to support dynamic access to multiple AWS services by using path-based routing. Which solution will meet these requirements?

A. Deploy an Application Load Balancer behind an Amazon API Gateway REST API. Configure IAM authorization.

B. Deploy an Application Load Balancer behind an Amazon API Gateway HTTP API. Use Amazon Cognito for authorization.

C. Deploy a Network Load Balancer behind an Amazon API Gateway REST API. Use an AWS Lambda function as a custom authorizer.

D. Deploy a Network Load Balancer behind an Amazon API Gateway HTTP API. Use Amazon Cognito for authorization.

Answer: C

QUESTION 1214

A company is performing a security review of its Amazon EMR API usage. The company's developers use an integrated development environment (IDE) that is hosted on Amazon EC2 instances. The IDE is configured to authenticate users to AWS by using access keys. Traffic between the company's EC2 instances and EMR cluster uses public IP addresses. A solutions architect needs to improve the company's overall security posture. The solutions architect needs to reduce the company's use of long-term credentials and to limit the amount of communication that uses public IP addresses. Which combination of steps will MOST improve the security of the company's architecture? (Choose two.)

A. Set up a gateway endpoint to the EMR cluster.

B. Set up interface VPC endpoints to connect to the EMR cluster.

C. Set up a private NAT gateway to connect to the EMR cluster.

D. Set up IAM roles for the developers to use to connect to the Amazon EMR API.

Answer: BD

QUESTION 1215

A company is implementing a new policy to enhance the security of its AWS environment. The policy requires all administrative actions that users perform on the AWS Management Console to be secured by multi-factor authentication (MFA). Which solution will allow the company to enforce this policy in the MOST operationally efficient way?

A. Enable MFA on the root account. Ensure that all administrators use the root account to perform administrative actions.

B. Create an IAM policy that requires MFA to be enabled for the IAM roles that administrators assume to perform administrative actions.

C. Configure an Amazon CloudWatch alarm that sends an email notification when an administrator performs an administrative action without MFA.

D. Use AWS Config to periodically audit IAM users and to automatically attach an IAM policy that requires MFA when AWS Config detects administrative actions.

Answer: B

QUESTION 1216

A company is building a new application that uses multiple serverless architecture components. The application architecture includes an Amazon API Gateway REST API and AWS Lambda functions to manage incoming requests. The company needs a service to send messages that the REST API receives to multiple target Lambda functions for processing. The service must filter messages so each target Lambda function receives only the messages the function needs. Which solution will meet these requirements with the LEAST operational overhead?

A. Send the requests from the REST API to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe multiple Amazon Simple Queue Service (Amazon SQS) queues to the SNS topic. Configure the target Lambda functions to poll the SQS queues.

B. Send the requests from the REST API to a set of Amazon EC2 instances that are configured to process messages. Configure the instances to filter messages and to invoke the target Lambda functions.

C. Send the requests from the REST API to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Configure Amazon MSK to publish the messages to the target Lambda functions.

D. Send the requests from the REST API to multiple Amazon Simple Queue Service (Amazon SQS) queues. Configure the target Lambda functions to poll the SQS queues.

Answer: C

QUESTION 1217

A solutions architect is investigating compute options for a critical analytics application. The application uses long-running processes to prepare and aggregate data. The processes cannot be interrupted. The application has a known baseline load. The application needs to handle occasional usage surges. Which solution will meet these requirements MOST cost-effectively?

A. Create an Amazon EC2 Auto Scaling group. Set the Min capacity and Desired capacity parameters to the number of instances required to handle the baseline load. Purchase Reserved Instances for the Auto Scaling group.

B. Create an Amazon EC2 Auto Scaling group.

Set the Min capacity, Max capacity, and Desired capacity parameters to the number of instances required to handle the baseline load. Use On-Demand Instances to address occasional usage surges.C. Create an Amazon EC2 Auto Scaling group. Set the Min capacity and Desired capacity parameters to the number of instances required to handle the baseline load. Purchase Reserved Instances for the Auto Scaling group. Use the OnDemandPercentageAboveBaseCapacity parameter to configure the launch template to launch Spot Instances.D. Re-architect the application to use AWS Lambda functions instead of Amazon EC2 instances. Purchase a one-year Compute Savings Plan to reduce the cost of Lambda usage. Answer: CQUESTION 1218A company wants to deploy an AWS Lambda function that will read and write objects to Amazon S3 bucket. The Lambda function must be connected to the company's VPC. The company must deploy the Lambda function only to private subnets in the VPC. The Lambda function must not be allowed to access the internet. Which solutions will meet these requirements? (Choose two.)A. Create a private NAT gateway to access the S3 bucket.B. Attach an Elastic IP address to the NAT gateway.C. Create a gateway VPC endpoint for the S3 bucket.D. Create an interface VPC endpoint for the S3 bucket.E. Create a public NAT gateway to access the S3 bucket. Answer: CDQUESTION 1219A company is creating a low-latency payment processing application that supports TLS connections from IPv4 clients. The application requires outbound access to the public internet. Users must access the application from a single entry point. The bank wants to use Amazon Elastic Container Service (Amazon ECS) tasks to deploy the application. The company wants to enable AWSVPC network mode. Which solution will meet these requirements MOST securely?A. Create a VPC that has an internet gateway, public subnets, and private subnets. Deploy a Network Load Balancer and a NAT gateway in the public subnets. Deploy the ECS tasks in the private subnets.B. Create a VPC that has an outbound-only internet gateway, public subnets, and private subnets. Deploy an Application Load Balancer and a NAT gateway in the public subnets. Deploy the ECS tasks in the private subnets.C. Create a VPC that has an internet gateway, public subnets, and private subnets. Deploy an Application Load Balancer in the public subnets. Deploy the ECS tasks in the public subnets.D. Create a VPC that has an outbound-only internet gateway, public subnets, and private subnets. Deploy a Network Load Balancer in the public subnets. Deploy the ECS tasks in the public subnets. Answer: AQUESTION 1220A company hosts an Amazon EC2 instance in a private subnet in a new VPC. The VPC also has a public subnet that has the default route set to an internet gateway. The private subnet does not have outbound internet access. The EC2 instance needs to have the ability to download monthly security updates from an outside vendor. However, the company must block any connections that are initiated from the internet. Which solution will meet these requirements?A. Configure the private subnet route table to use the internet gateway as the default route.B. Create a NAT gateway in the public subnet. Configure the private subnet route table to use the NAT gateway as the default route.C. Create a NAT instance in the private subnet. Configure the private subnet route table to use the NAT instance as the default route.D. Create a NAT instance in the private subnet. Configure the private subnet route table to use the internet gateway as the default route. Answer: BQUESTION 1221A company wants to create a payment processing application. The application must run when a payment record arrives in an existing Amazon S3 bucket. The application must process each payment record exactly once. The company wants to use an AWS Lambda function to process the payments. Which solution will meet these requirements?A. Configure the existing S3 bucket to send object creation events to Amazon EventBridge. Configure EventBridge to route events to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Configure the Lambda function to run when a new event arrives in the SQS queue.B. Configure the existing S3 bucket to send object creation events to an Amazon Simple Notification Service (Amazon SNS) topic. Configure the Lambda function to run when a new event arrives in the SNS topic.C. Configure the existing S3 bucket to send object creation events to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the Lambda function to run when a new event arrives in the SQS queue.D. Configure the existing S3 bucket to send object creation events directly to the Lambda function. Configure the Lambda function to handle object creation events and to process the payments. Answer: BQUESTION 1222A company is redesigning a static website. The company needs a solution to host the new website in the company's AWS account. The solution must be secure and scalable. Which combination of solutions will meet these requirements? (Choose three.)A. Configure an Amazon CloudFront distribution. Set the Amazon S3 bucket as the origin.B. Associate an AWS Certificate Manager (ACM) TLS certificate to the Amazon CloudFront distribution.C. Enable static website hosting for the Amazon S3 bucket.D. Create an Amazon S3 bucket to store the static website content.E. Export the website's SSL/TLS certificate from AWS Certificate Manager (ACM) to the root of the Amazon S3 bucket.F. Turn off Block Public Access for the Amazon S3 bucket. Answer: ABDQUESTION 1223A global company is migrating its workloads from an on-premises data center to AWS. The AWS environment includes multiple AWS accounts. IAM roles, AWS Config rules, and a VPC. The company wants an automated process to provision new accounts on demand when the company's business units require new accounts. Which solution will meet these requirements with LEAST effort?A. Use AWS Control Tower to set up an organization in AWS Organizations. Use AWS Control Tower Account Factory for Terraform (AFT) to provision new AWS accounts.B. Create an organization in AWS Organizations. Use the AWS CLI

CreateAccount API action to provision new AWS accounts. Organize the business units with organizational units (OUs).C. Create an AWS Lambda function that uses the AWS Organizations API to create new accounts. Invoke the Lambda function from an AWS CloudFormation template in AWS Service Catalog.D. Create an organization in AWS Organizations. Use AWS Step Functions to orchestrate the account creation process. Send account creation requests to an Amazon API Gateway API endpoint to invoke an AWS Lambda function that creates new accounts.

Answer: AQUESTION 1224A company wants to design a microservices architecture for an application. Each microservice must perform operations that can be completed within 30 seconds. The microservices need to expose RESTful APIs and must automatically scale in response to varying loads. The APIs must also provide client access control and rate limiting to maintain equitable usage and service availability. Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 to host each microservice. Use Amazon API Gateway to manage the RESTful API requests.

B. Deploy each microservice as a set of AWS Lambda functions. Use Amazon API Gateway to manage the RESTful API requests.

C. Host each microservice on Amazon EC2 instances in Auto Scaling groups behind an Elastic Load Balancing (ELB) load balancer. Use the ELB to manage the RESTful API requests.

D. Deploy each microservice on Amazon Elastic Beanstalk. Use Amazon CloudFront to manage the RESTful API requests.

Answer: CQUESTION 1225A company is launching a new gaming application. The company will use Amazon EC2 Auto Scaling groups to deploy the application. The application stores user data in a relational database. The company has office locations around the world that need to run analytics on the user data in the database. The company needs a cost-effective database solution that provides cross-Region disaster recovery with low-latency read performance across AWS Regions. Which solution will meet these requirements?

A. Create an Amazon ElastiCache for Redis cluster in the Region where the application is deployed. Create read replicas in Regions where the company offices are located. Ensure the company offices read from the read replica instances.

B. Create Amazon DynamoDB global tables. Deploy the tables to the Regions where the company offices are located and to the Region where the application is deployed. Ensure that each company office reads from the tables that are in the same Region as the office.

C. Create an Amazon Aurora global database. Configure the primary cluster to be in the Region where the application is deployed. Configure the secondary Aurora replicas to be in the Regions where the company offices are located. Ensure the company offices read from the Aurora replicas.

D. Create an Amazon RDS Multi-AZ DB cluster deployment in the Region where the application is deployed. Ensure the company offices read from read replica instances.

Answer: AQUESTION 1226An international company needs to share data from an Amazon S3 bucket to employees who are located around the world. The company needs a secure solution to provide employees with access to the S3 bucket. The employees are already enrolled in AWS IAM Identity Center. Which solution will meet these requirements with the LEAST operational overhead?

A. Create a help desk application to generate an Amazon S3 presigned URL for each employee. Configure the presigned URLs to have short expirations. Instruct employees to contact the company help desk to receive a presigned URL to access the S3 bucket.

B. Create a group for Amazon S3 access in IAM Identity Center. Add the employees who require access to the S3 bucket to the group. Create an IAM policy to allow Amazon S3 access from the group. Instruct employees to use the AWS access portal to access the AWS Management Console and navigate to the S3 bucket.

C. Create an Amazon S3 File Gateway. Create one share for data uploads and a second share for data downloads. Set up an SFTP service on an Amazon EC2 instance. Mount the shares to the EC2 instance. Instruct employees to use the SFTP server.

D. Configure AWS Transfer Family SFTP endpoints. Select the custom identity provider option. Use AWS Secrets Manager to manage the user credentials. Instruct employees to use Transfer Family SFTP.

Answer: CQUESTION 1227A company hosts a website analytics application on a single Amazon EC2 On-Demand Instance. The analytics application is highly resilient and is designed to run in stateless mode. The company notices that the application is showing signs of performance degradation during busy times and is presenting 5xx errors. The company needs to make the application scale seamlessly. Which solution will meet these requirements MOST cost-effectively?

A. Create an Amazon Machine Image (AMI) of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use an Application Load Balancer to distribute the load across the two EC2 instances.

B. Create an Amazon Machine Image (AMI) of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.

C. Create an AWS Lambda function to stop the EC2 instance and change the instance type. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization is more than 75%.

D. Create an Amazon Machine Image (AMI) of the web application. Apply the AMI to a launch template. Create an Auto Scaling group that includes the launch template. Configure the launch template to use a Spot Fleet. Attach an Application Load Balancer to the Auto Scaling group.

Answer: DExplanation: Auto Scaling is the most effective solution for ensuring seamless scalability of a stateless application. Key points: D leverages Auto Scaling with a Spot Fleet for cost efficiency and attaches an ALB to distribute traffic. A and B do not provide automated scaling and would require manual intervention to add more instances. C changes the instance type but does not scale out horizontally, which is required.

here.QUESTION 1228A company hosts its applications in multiple private and public subnets in a VPC. The applications in the private subnets need to access an API. The API is available on the internet and is hosted in the company's on-premises data center. A solutions architect needs to establish connectivity for applications in the private subnets. Which solution will meet these requirements MOST cost-effectively?A. Create a transit gateway to connect the VPC to the on-premises network. Use the transit gateway to route API calls from the private subnets to the on-premises data center.B. Create a NAT gateway in the public subnet of the VPC. Use the NAT gateway to allow the private subnets to access the API over the internet.C. Establish an AWS PrivateLink connection to connect the VPC to the on-premises network. Use PrivateLink to make API calls from the private subnets to the on-premises data center.D. Implement an AWS Site-to-Site VPN connection between the VPC and the on-premises data center. Use the VPN connection to make API calls from the private subnets to the on-premises data center.Answer: DExplanation: AWS Site-to-Site VPN is a cost-effective way to securely connect your on-premises data center with AWS resources. In this scenario: Applications in private subnets require access to the API hosted in the on-premises data center. A Site-to-Site VPN connection is a secure and cost-efficient option to route traffic between the VPC and on-premises resources. Transit Gateway and PrivateLink are not cost-effective for this use case. NAT Gateway only provides internet access for private subnets, which is not suitable for reaching an on-premises resource.QUESTION 1229A healthcare company is developing an AWS Lambda function that publishes notifications to an encrypted Amazon Simple Notification Service (Amazon SNS) topic. The notifications contain protected health information (PHI). The SNS topic uses AWS Key Management Service (AWS KMS) customer-managed keys for encryption. The company must ensure that the application has the necessary permissions to publish messages securely to the SNS topic. Which combination of steps will meet these requirements? (Choose three.)A. Create a resource policy for the SNS topic that allows the Lambda function to publish messages to the topic.B. Use server-side encryption with AWS KMS keys (SSE-KMS) for the SNS topic instead of customer-managed keys.C. Create a resource policy for the encryption key that the SNS topic uses that has the necessary AWS KMS permissions.D. Specify the Lambda function's Amazon Resource Name (ARN) in the SNS topic's resource policy.E. Associate an Amazon API Gateway HTTP API with the SNS topic to control access to the topic by using API Gateway resource policies.F. Configure a Lambda execution role that has the necessary IAM permissions to use a customer-managed key in AWS KMS.Answer: ACFExplanation: To securely publish messages to an encrypted Amazon SNS topic, the following steps are required:A. Resource policy for SNS topic: Ensures that the Lambda function is explicitly allowed to publish messages to the topic.C. Resource policy for KMS key: Provides the necessary permissions to use the customer-managed key for encryption.F. Lambda execution role: Grants the Lambda function the necessary IAM permissions to use the encryption key.QUESTION 1230A company stores sensitive customer data in an Amazon DynamoDB table. The company frequently updates the data. The company wants to use the data to personalize offers for customers. The company's analytics team has its own AWS account. The analytics team runs an application on Amazon EC2 instances that needs to process data from the DynamoDB tables. The company needs to follow security best practices to create a process to regularly share data from DynamoDB to the analytics team. Which solution will meet these requirements?A. Export the required data from the DynamoDB table to an Amazon S3 bucket as multiple JSON files. Provide the analytics team with the necessary IAM permissions to access the S3 bucket.B. Allow public access to the DynamoDB table. Create an IAM user that has permission to access DynamoDB. Share the IAM user with the analytics team.C. Allow public access to the DynamoDB table. Create an IAM user that has read-only permission for DynamoDB. Share the IAM user with the analytics team.D. Create a cross-account IAM role. Create an IAM policy that allows the AWS account ID of the analytics team to access the DynamoDB table. Attach the IAM policy to the IAM role. Establish a trust relationship between accounts.Answer: DExplanation: Using cross-account IAM roles is the most secure and scalable way to share data between AWS accounts. A trust relationship allows the analytics team's account to assume the role in the main account and access the DynamoDB table directly. A is feasible but involves data duplication and additional costs for storing the JSON files in S3. B and C violate security best practices by allowing public access to sensitive data and sharing credentials, which is highly discouraged.QUESTION 1231A financial service company has a two-tier consumer banking application. The frontend serves static web content. The backend consists of APIs. The company needs to migrate the frontend component to AWS. The backend of the application will remain on-premises. The company must protect the application from common web vulnerabilities and attacks. Which solution will meet these requirements with the LEAST operational overhead?A. Migrate the frontend to Amazon EC2 instances. Deploy an Application Load Balancer (ALB) in front of the instances. Use the instances to invoke the on-premises APIs. Associate AWS WAF rules with the instances.B. Deploy the frontend as an Amazon CloudFront distribution that has multiple origins. Configure one origin to be an Amazon S3 bucket that serves the static web content. Configure a second origin to route traffic to the on-premises APIs based on the URL pattern. Associate AWS WAF rules with the distribution.C. Migrate the frontend to Amazon EC2 instances. Deploy a Network Load Balancer (NLB) in front of the instances. Use the instances to invoke the

on-premises APIs. Create an AWS Network Firewall instance. Route all traffic through the Network Firewall instance.D. Deploy the frontend as a static website based on an Amazon S3 bucket. Use an Amazon API Gateway REST API and a set of Amazon EC2 instances to invoke the on-premises APIs. Associate AWS WAF rules with the REST API and the S3 bucket. Answer: BExplanation: Comprehensive Deploying the frontend as a CloudFront distribution with multiple origins provides an efficient and scalable solution. Using WAF rules with CloudFront protects against web vulnerabilities, while the multi-origin configuration allows traffic routing to the on-premises backend APIs. This approach minimizes operational overhead compared to managing EC2 instances. QUESTION 1232A company is deploying a new application to a VPC on existing Amazon EC2 instances. The application has a presentation tier that uses an Auto Scaling group of EC2 instances. The application also has a database tier that uses an Amazon RDS Multi-AZ database. The VPC has two public subnets that are split between two Availability Zones. A solutions architect adds one private subnet to each Availability Zone for the RDS database. The solutions architect wants to restrict network access to the RDS database to block access from EC2 instances that do not host the new application. Which solution will meet this requirement?A. Modify the RDS database security group to allow traffic from a CIDR range that includes IP addresses of the EC2 instances that host the new application.B. Associate a new ACL with the private subnets. Deny all incoming traffic from IP addresses that belong to any EC2 instance that does not host the new application.C. Modify the RDS database security group to allow traffic from the security group that is associated with the EC2 instances that host the new application.D. Associate a new ACL with the private subnets. Deny all incoming traffic except for traffic from a CIDR range that includes IP addresses of the EC2 instances that host the new application. Answer: CExplanation: AWS Security Groups: Security groups operate at the instance level, making them the ideal tool for controlling access to specific resources such as an Amazon RDS database. By default, security groups deny all incoming traffic. You can allow access by explicitly specifying another security group. Associating an RDS database security group with the EC2 instances' security group ensures only the specified EC2 instances can access the RDS database. QUESTION 1233A developer is creating a serverless application that performs video encoding. The encoding process runs as background jobs and takes several minutes to encode each video. The process must not send an immediate result to users. The developer is using Amazon API Gateway to manage an API for the application. The developer needs to run test invocations and request validations. The developer must distribute API keys to control access to the API. Which solution will meet these requirements?A. Create an HTTP API. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the HTTP API. Use the Event invocation type to call the Lambda function.B. Create a REST API with the default endpoint type. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the REST API. Use the Event invocation type to call the Lambda function.C. Create an HTTP API. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the HTTP API. Use the RequestResponse invocation type to call the Lambda function.D. Create a REST API with the default endpoint type. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the REST API. Use the RequestResponse invocation type to call the Lambda function. Answer: BExplanation: The Event invocation type is asynchronous, meaning the Lambda function does not send an immediate result to the API Gateway and processes the request in the background. This is ideal for video encoding tasks that take time. REST API vs. HTTP API: REST APIs support advanced features like API keys, request validation, and throttling that HTTP APIs do not support fully. Since the developer needs API keys and request validations, a REST API is the correct choice. Integration with Lambda: AWS Lambda integration is seamless with REST APIs, and using the Event invocation ensures asynchronous processing. QUESTION 1234A developer needs to export the contents of several Amazon DynamoDB tables into Amazon S3 buckets to comply with company data regulations. The developer uses the AWS CLI to run commands to export from each table to the proper S3 bucket. The developer sets up AWS credentials correctly and grants resources appropriate permissions. However, the exports of some tables fail. What should the developer do to resolve this issue?A. Ensure that point-in-time recovery is enabled on the DynamoDB tables.B. Ensure that the target S3 bucket is in the same AWS Region as the DynamoDB table.C. Ensure that DynamoDB streaming is enabled for the tables.D. Ensure that DynamoDB Accelerator (DAX) is enabled. Answer: AExplanation: To export data from DynamoDB to Amazon S3, point-in-time recovery (PITR) must be enabled for the tables. This feature creates a snapshot of the data, which is essential for exports. QUESTION 1235A developer is creating an ecommerce workflow in an AWS Step Functions state machine that includes an HTTP Task state. The task passes shipping information and order details to an endpoint. The developer needs to test the workflow to confirm that the HTTP headers and body are correct and that the responses meet expectations. Which solution will meet these requirements?A. Use the TestState API to invoke only the HTTP Task. Set the inspection level to TRACE.B. Use the TestState API to invoke the state machine. Set the inspection level to DEBUG.C. Use the data flow simulator to invoke only the HTTP Task. View the request and response data.D. Change the log level of the state machine to ALL. Run the state machine. Answer: DExplanation: Changing the log level to ALL enables capturing detailed request and response data. This helps verify HTTP headers,

body, and responses.QUESTION 1236A developer used the AWS SDK to create an application that aggregates and produces log records for 10 services. The application delivers data to an Amazon Kinesis Data Streams stream. Each record contains a log message with a service name, creation timestamp, and other log information. The stream has 15 shards in provisioned capacity mode. The stream uses service name as the partition key. The developer notices that when all the services are producing logs, ProvisionedThroughputExceededException errors occur during PutRecord requests. The stream metrics show that the write capacity the applications use is below the provisioned capacity. How should the developer resolve this issue?A. Change the capacity mode from provisioned to on-demand.B. Double the number of shards until the throttling errors stop occurring.C. Change the partition key from service name to creation timestamp.D. Use a separate Kinesis stream for each service to generate the logs. Answer: CExplanation: Using "service name" as the partition key results in uneven data distribution. Some shards may become hot due to excessive logs from certain services, leading to throttling errors. Changing the partition key to "creation timestamp" ensures a more even distribution of records across shards.QUESTION 1237A company has a serverless web application that is comprised of AWS Lambda functions. The application experiences spikes in traffic that cause increased latency because of cold starts. The company wants to improve the application's ability to handle traffic spikes and to minimize latency. The solution must optimize costs during periods when traffic is low. Which solution will meet these requirements?A. Configure provisioned concurrency for the Lambda functions. Use AWS Application Auto Scaling to adjust the provisioned concurrency.B. Launch Amazon EC2 instances in an Auto Scaling group. Add a scheduled scaling policy to launch additional EC2 instances during peak traffic periods.C. Configure provisioned concurrency for the Lambda functions. Set a fixed concurrency level to handle the maximum expected traffic.D. Create a recurring schedule in Amazon EventBridge Scheduler. Use the schedule to invoke the Lambda functions periodically to warm the functions. Answer: AExplanation: AWS Lambda's provisioned concurrency ensures that a predefined number of execution environments are pre-warmed and ready to handle requests, reducing latency during traffic spikes. This solution optimizes costs during low-traffic periods when combined with AWS Application Auto Scaling to dynamically adjust the provisioned concurrency based on demand.QUESTION 1238A solutions architect is designing the storage architecture for a new web application used for storing and viewing engineering drawings. All application components will be deployed on the AWS infrastructure. The application design must support caching to minimize the amount of time that users wait for the engineering drawings to load. The application must be able to store petabytes of data. Which combination of storage and caching should the solutions architect use?A. Amazon S3 with Amazon CloudFrontB. Amazon S3 Glacier Deep Archive with Amazon ElastiCacheC. Amazon Elastic Block Store (Amazon EBS) volumes with Amazon CloudFrontD. AWS Storage Gateway with Amazon ElastiCache Answer: AExplanation: Amazon S3 provides highly scalable and durable storage for petabytes of data. Amazon CloudFront, as a content delivery network (CDN), caches frequently accessed data at edge locations to reduce latency. This combination is ideal for storing and accessing engineering drawings.QUESTION 1239A media company has an ecommerce website to sell music. Each music file is stored as an MP3 file. Premium users of the website purchase music files and download the files. The company wants to store music files on AWS. The company wants to provide access only to the premium users. The company wants to use the same URL for all premium users. Which solution will meet these requirements?A. Store the MP3 files on a set of Amazon EC2 instances that have Amazon Elastic Block Store (Amazon EBS) volumes attached. Manage access to the files by creating an IAM user and an IAM policy for each premium user.B. Store all the MP3 files in an Amazon S3 bucket. Create a presigned URL for each MP3 file. Share the presigned URLs with the premium users.C. Store all the MP3 files in an Amazon S3 bucket. Create an Amazon CloudFront distribution that uses the S3 bucket as the origin. Generate CloudFront signed cookies for the music files. Share the signed cookies with the premium users.D. Store all the MP3 files in an Amazon S3 bucket. Create an Amazon CloudFront distribution that uses the S3 bucket as the origin. Use a CloudFront signed URL for each music file. Share the signed URLs with the premium users. Answer: CExplanation: CloudFront signed cookies allow the company to provide access to premium users while maintaining a single, consistent URL. This approach is simpler and more scalable than managing presigned URLs for each file.QUESTION 1240A company is building a serverless application to process orders from an e-commerce site. The application needs to handle bursts of traffic during peak usage hours and to maintain high availability. The orders must be processed asynchronously in the order the application receives them. Which solution will meet these requirements?A. Use an Amazon Simple Notification Service (Amazon SNS) topic to receive orders. Use an AWS Lambda function to process the orders.B. Use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to receive orders. Use an AWS Lambda function to process the orders.C. Use an Amazon Simple Queue Service (Amazon SQS) standard queue to receive orders. Use AWS Batch jobs to process the orders.D. Use an Amazon Simple Notification Service (Amazon SNS) topic to receive orders. Use AWS Batch jobs to process the orders. Answer: BQUESTION 1241An e-commerce company has an application that uses Amazon DynamoDB tables configured with provisioned capacity. Order data is stored in a table named Orders. The Orders table has a primary key of order-ID and a sort

key of product-ID. The company configured an AWS Lambda function to receive DynamoDB streams from the Orders table and update a table named Inventory. The company has noticed that during peak sales periods, updates to the Inventory table take longer than the company can tolerate. Which solutions will resolve the slow table updates? (Choose two.)  
A. Add a global secondary index to the Orders table. Include the product-ID attribute.  
B. Set the batch size attribute of the DynamoDB streams to be based on the size of items in the Orders table.  
C. Increase the DynamoDB table provisioned capacity by 1,000 write capacity units (WCUs).  
D. Increase the DynamoDB table provisioned capacity by 1,000 read capacity units (RCUs).  
E. Increase the timeout of the Lambda function to 15 minutes.

Answer: BCQUESTION 1242

A company has an e-commerce site. The site is designed as a distributed web application hosted in multiple AWS accounts under one AWS Organizations organization. The web application is comprised of multiple microservices. All microservices expose their AWS services either through Amazon CloudFront distributions or public Application Load Balancers (ALBs). The company wants to protect public endpoints from malicious attacks and monitor security configurations. Which solution will meet these requirements with the LEAST operational overhead?  
A. Use AWS WAF to protect the public endpoints. Use AWS Firewall Manager from a dedicated security account to manage rules in AWS WAF. Use AWS Config rules to monitor the Regional and global WAF configurations.  
B. Use AWS WAF to protect the public endpoints. Apply AWS WAF rules in each account. Use AWS Config rules and AWS Security Hub to monitor the WAF configurations of the ALBs and the CloudFront distributions.  
C. Use AWS WAF to protect the public endpoints. Use AWS Firewall Manager from a dedicated security account to manage the rules in AWS WAF. Use Amazon Inspector and AWS Security Hub to monitor the WAF configurations of the ALBs and the CloudFront distributions.  
D. Use AWS Shield Advanced to protect the public endpoints. Use AWS Config rules to monitor the Shield Advanced configuration for each account.

Answer: AQUESTION 1243

A company has a serverless web application that is comprised of AWS Lambda functions. The application experiences spikes in traffic that cause increased latency because of cold starts. The company wants to improve the application's ability to handle traffic spikes and to minimize latency. The solution must optimize costs during periods when traffic is low. Which solution will meet these requirements?  
A. Configure provisioned concurrency for the Lambda functions. Use AWS Application Auto Scaling to adjust the provisioned concurrency.  
B. Launch Amazon EC2 instances in an Auto Scaling group. Add a scheduled scaling policy to launch additional EC2 instances during peak traffic periods.  
C. Configure provisioned concurrency for the Lambda functions. Set a fixed concurrency level to handle the maximum expected traffic.  
D. Create a recurring schedule in Amazon EventBridge Scheduler. Use the schedule to invoke the Lambda functions periodically to warm the functions.

Answer: AQUESTION 1244

A financial services company has a two-tier consumer banking application. The frontend serves static web content. The backend consists of APIs. The company needs to migrate the frontend component to AWS. The backend of the application will remain on-premises. The company must protect the application from common web vulnerabilities and attacks. Which solution will meet these requirements?  
A. Migrate the frontend to Amazon EC2 instances. Deploy an Application Load Balancer (ALB) in front of the instances. Use the instances to invoke the on-premises APIs. Associate AWS WAF rules with the instances.  
B. Deploy the frontend as an Amazon CloudFront distribution that has multiple origins. Configure one origin to be an Amazon S3 bucket that serves the static web content. Configure a second origin to route traffic to the on-premises APIs based on the URL pattern. Associate AWS WAF rules with the distribution.  
C. Migrate the frontend to Amazon EC2 instances. Deploy a Network Load Balancer (NLB) in front of the instances. Use the instances to invoke the on-premises APIs. Create an AWS Network Firewall instance. Route all traffic through the Network Firewall instance.  
D. Deploy the frontend as a static website based on an Amazon S3 bucket. Use an Amazon API Gateway REST API and a set of Amazon EC2 instances to invoke the on-premises APIs. Associate AWS WAF rules with the REST API and the S3 bucket.

Answer: BQUESTION 1245

A company is deploying a critical application by using Amazon RDS for MySQL. The application must be highly available and must recover automatically. The company needs to support interactive users (transactional queries) and batch reporting (analytical queries) with no more than a 4-hour lag. The analytical queries must not affect the performance of the transactional queries.  
A. Configure Amazon RDS for MySQL in a Multi-AZ DB instance deployment with one standby instance. Point the transactional queries to the primary DB instance. Point the analytical queries to a secondary DB instance that runs in a different Availability Zone.  
B. Configure Amazon RDS for MySQL in a Multi-AZ DB cluster deployment with two standby instances. Point the transactional queries to the primary DB instance. Point the analytical queries to the reader endpoint.  
C. Configure Amazon RDS for MySQL to use multiple read replicas across multiple Availability Zones. Point the transactional queries to the primary DB instance. Point the analytical queries to one of the replicas in a different Availability Zone.  
D. Configure Amazon RDS for MySQL as the primary database for the transactional queries with automated backups enabled. Configure automated backups. Each night, create a read-only database from the most recent snapshot to support the analytical queries. Terminate the previously created database.

Answer: CQUESTION 1246

A company plans to use an Amazon S3 bucket to archive backup data. Regulations require the company to retain the backup data for 7 years. During the retention period, the company must

prevent users, including administrators, from deleting the data. The company can delete the data after 7 years. Which solution will meet these requirements?

- A. Create an S3 bucket policy that denies delete operations for 7 years. Create an S3 Lifecycle policy to delete the data after 7 years.
- B. Create an S3 Object Lock default retention policy that retains data for 7 years in governance mode. Create an S3 Lifecycle policy to delete the data after 7 years.
- C. Create an S3 Object Lock default retention policy that retains data for 7 years in compliance mode. Create an S3 Lifecycle policy to delete the data after 7 years.
- D. Create an S3 Batch Operations job to set a legal hold on each object for 7 years. Create an S3 Lifecycle policy to delete the data after 7 years.

Answer: C

QUESTION 1247

A company runs multiple applications on Amazon EC2 instances in a VPC. Application A runs in a private subnet that has a custom route table and network ACL. Application B runs in a second private subnet in the same VPC. The company needs to prevent Application A from sending traffic to Application B. Which solution will meet this requirement?

- A. Add an deny outbound rule to a security group associated with Application B. Configure the rule to prevent Application B from sending traffic to Application A.
- B. Add an deny outbound rule to a security group associated with Application A. Configure the rule to prevent Application A from sending traffic to Application B.
- C. Add an deny outbound rule to the custom network ACL for the Application B subnet. Configure the rule to prevent Application B from sending traffic to the IP addresses associated with Application A.
- D. Add an deny outbound rule to the custom network ACL for the Application A subnet. Configure the rule to prevent Application A from sending traffic to the IP addresses associated with Application B.

Answer: D

QUESTION 1248

A company generates approximately 20 GB of data multiple times each day. The company uses AWS DataSync to copy all data from on-premises storage to Amazon S3 every 6 hours for further processing. The analytics team wants to modify the copy process to copy only data relevant to the analytics team and ignore the rest of the data. The team wants to copy data as soon as possible and receive a notification when the copy process is finished. Which combination of steps will meet these requirements MOST cost-effectively? (Choose three.)

- A. Modify the data generation process on-premises to create a manifest file at the end of the copy process with the names of the objects to be copied to Amazon S3. Create a custom script to upload the manifest file to an S3 bucket.
- B. Modify the data generation process on-premises to create a manifest file at the end of the copy process with the names of the objects to be copied to Amazon S3. Create an AWS Lambda function to load the manifest file data into an Amazon DynamoDB table.
- C. Create an AWS Lambda function that Amazon EventBridge invokes when the manifest file is loaded into Amazon DynamoDB. Configure the Lambda function to copy the data from on-premises storage to the S3 bucket that uses the manifest file.
- D. Create an AWS Lambda function that an S3 Event Notification invokes when the manifest file is uploaded. Configure the Lambda function to invoke the DataSync task by calling the StartTaskExecution API action with a manifest.
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic. Create an Amazon EventBridge rule to send an email notification to the SNS topic when the DataSync task execution status changes to SUCCESS or to ERROR.
- F. Create an Amazon Simple Notification Service (Amazon SNS) topic. Create an AWS Lambda function to send an email notification to the SNS topic when the DataSync task execution status changes to SUCCESS or to ERROR.

Answer: ADE

QUESTION 1249

A company uses an AWS Transfer for SFTP public server endpoint and Amazon S3 storage to host large datasets for its customers. The company provides customers SSH private keys to authenticate and download their datasets. The Transfer for SFTP server is configured with structured logging that is saved to an S3 bucket. The company wants to charge customers based on their monthly data download usage. Which solution will meet these requirements?

- A. Configure VPC Flow Logs to write to a new S3 bucket. Run monthly queries on the flow logs to identify customer usage and calculate cost. Add the charges to the customers' monthly bills.
- B. Each month, use AWS Cost Explorer to examine the costs for Transfer for SFTP and obtain a breakdown by customer. Add the charges to the customers' monthly bills.
- C. Enable requester pays on the S3 bucket that hosts the software. Allocate the charges to each customer based on the customer's requests.
- D. Run Amazon Athena queries on the logging S3 bucket monthly to identify customer usage and calculate costs. Add the charges to the customers' monthly bills.

Answer: D

QUESTION 1250

A company collects data for temperature, humidity, and atmospheric pressure in cities across multiple continents. The average volume of data that the company collects from each site daily is 500 GB. Each site has a high-speed internet connection. The company wants to aggregate the data from all these global sites as quickly as possible in a single Amazon S3 bucket. The solution must minimize operational complexity. Which solution meets these requirements?

- A. Turn on S3 Transfer Acceleration on the destination S3 bucket. Use multipart uploads to directly upload site data to the destination S3 bucket.
- B. Upload the data from each site to an S3 bucket in the closest Region. Use S3 Cross-Region Replication to copy objects to the destination S3 bucket. Then remove the data from the origin S3 bucket.
- C. Schedule AWS Snowball Edge Storage Optimized device jobs daily to transfer data from each site to the closest Region. Use S3 Cross-Region Replication to copy objects to the destination S3 bucket.
- D. Upload the data from each site to an Amazon EC2 instance in the closest Region. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. At regular intervals, take an EBS snapshot and copy it to the Region that contains the destination S3 bucket. Restore the EBS volume in that Region.

Answer: A

QUESTION 1251

An e-commerce company

stores inventory, order, and user information in multiple Amazon Redshift clusters. The Redshift clusters must comply with the company's security policies. The company must receive notifications about any security configuration violations. Which solution will meet these requirements?A. Create an Amazon EventBridge rule that uses the Redshift clusters as the source. Create an AWS Lambda function to evaluate the Redshift cluster security configuration. Configure the Lambda function to notify the company of any violations of the security policies. Add the Lambda function as a target of the EventBridge rule.B. Create an AWS Lambda function to check the validity of the Redshift cluster security configurations. Create an Amazon EventBridge rule that invokes the Lambda function when Redshift clusters are created. Notify the company of any violations of security policies.C. Set up Amazon Redshift Advisor in the company's AWS account to monitor cluster configurations. Configure Redshift Advisor to generate notifications for security items that the company must address.D. Create an AWS Lambda function to check the Redshift clusters for any violation of the security configurations. Create an AWS Config custom rule to invoke the Lambda function when Redshift cluster security configurations are modified. Provide the compliance state of each Redshift cluster to AWS Config. Configure AWS Config to notify the company of any violations of the security policies.

Answer: D

QUESTION 1252

A company uses AWS Cost Explorer to monitor its AWS costs. The company notices that Amazon Elastic Block Store (Amazon EBS) storage and snapshot costs increase every month. However, the company does not purchase additional EBS storage every month. The company wants to optimize monthly costs for its current storage usage. Which solution will meet these requirements with the LEAST operational overhead?A. Use logs in Amazon CloudWatch Logs to monitor the storage utilization of Amazon EBS. Use Amazon EBS Elastic Volumes to reduce the size of the EBS volumes.B. Use a custom script to monitor space usage. Use Amazon EBS Elastic Volumes to reduce the size of the EBS volumes.C. Delete all expired and unused snapshots to reduce snapshot costs.D. Delete all nonessential snapshots. Use Amazon Data Lifecycle Manager to create and manage the snapshots according to the company's snapshot policy requirements.

Answer: D

Explanation: Amazon Data Lifecycle Manager (DLM) automates the creation, retention, and deletion of EBS snapshots. This allows organizations to define policies that ensure snapshots are only kept as long as needed, reducing costs automatically and minimizing manual effort. AWS recommends using DLM for optimizing storage and managing backup lifecycle with minimal overhead.

QUESTION 1253

A company uses Amazon Elastic Container Service (Amazon ECS) to run workloads that belong to service teams. Each service team uses an owner tag to specify the ECS containers that the team owns. The company wants to generate an AWS Cost Explorer report that shows how much each service team spends on ECS containers on a monthly basis. Which combination of steps will meet these requirements in the MOST operationally efficient way? (Choose two.)

A. Create a custom report in Cost Explorer. Apply a filter for Amazon ECS.B. Create a custom report in Cost Explorer. Apply a filter for the owner resource tag.C. Set up AWS Compute Optimizer. Review the rightsizing recommendations.D. Activate the owner tag as a cost allocation tag. Group the Cost Explorer report by linked account.

Answer: A, C

Explanation: To allocate costs based on team ownership, AWS recommends tagging resources using cost allocation tags. Activating the "owner" tag as a cost allocation tag allows AWS Cost Explorer to categorize and group spending. Additionally, filtering for Amazon ECS services enables the visibility into service-specific usage. This approach is both scalable and operationally efficient.

QUESTION 1254

A company is planning to run an AI/ML workload on AWS. The company needs to train a model on a dataset that is in Amazon S3 Standard. A model training application requires multiple compute nodes and single-digit millisecond access to the data. Which solution will meet these requirements in the MOST cost-effective way?A. Move the data to S3 Intelligent-Tiering. Point the model training application to S3 Intelligent-Tiering as the data source.B. Add partitions to the S3 bucket by adding random prefixes. Reconfigure the model training application to point to the new prefixes as the data source.C. Move the data to S3 Express One Zone. Point the model training application to S3 Express One Zone as the data source.D. Move the data to a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume attached to an Amazon EC2 instance. Point the model training application to the gp3 volume as the data source.

Answer: C

Explanation: Amazon S3 Express One Zone provides single-digit millisecond latency and high throughput, making it ideal for ML workloads that require multiple compute nodes and fast access. It is also more cost-effective than traditional file or block storage for temporary, high-speed needs.

QUESTION 1255

A company is building a mobile gaming app. The company wants to serve users from around the world with low latency. The company needs a scalable solution to host the application and to route user requests to the location that is nearest to each user. Which solution will meet these requirements?A. Use an Application Load Balancer to route requests to Amazon EC2 instances that are deployed across multiple Availability Zones.B. Use a Regional Amazon API Gateway REST API to route requests to AWS Lambda functions.C. Use an edge-optimized Amazon API Gateway REST API to route requests to AWS Lambda functions.D. Use an Application Load Balancer to route requests to containers in an Amazon ECS cluster.

Answer: C

Explanation: Edge-optimized API Gateway endpoints utilize the Amazon CloudFront global network to decrease latency for clients globally. This setup ensures that the request is routed to the closest edge location, significantly

reducing response time and improving performance for worldwide users.

**QUESTION 1256** A solutions architect is designing a web application that will run on Amazon EC2 instances behind an Application Load Balancer (ALB). The company strictly requires that the application be resilient against malicious internet activity and attacks, and protect against new common vulnerabilities and exposures. What should the solutions architect recommend?

A. Leverage Amazon CloudFront with the ALB endpoint as the origin.

B. Deploy an appropriate managed rule for AWS WAF and associate it with the ALB.

C. Subscribe to AWS Shield Advanced and ensure common vulnerabilities and exposures are blocked.

D. Configure network ACLs and security groups to allow only ports 80 and 443 to access the EC2 instances.

**Answer: B**

**Explanation:** AWS WAF allows web applications to protect themselves from common web exploits and vulnerabilities. Using AWS managed rule groups ensures protection against known attack patterns, such as SQL injection and cross-site scripting. Associating AWS WAF with the ALB provides application-layer security and real-time threat mitigation.

**QUESTION 1257** A company has an organization in AWS Organizations that has all features enabled. The company has multiple Amazon S3 buckets in multiple AWS Regions around the world. The S3 buckets contain sensitive data. The company needs to ensure that no personally identifiable information (PII) is stored in the S3 buckets. The company also needs a scalable solution to identify PII. Which solution will meet these requirements?

A. In the Organizations management account, configure an Amazon Macie administrator IAM user as the delegated administrator for the global organization. Use the Macie administrator user to configure Macie settings to scan for PII.

B. For each Region in the Organizations management account, designate a delegated Amazon Macie administrator account. In the Macie administrator account, add all accounts in the organization. Use the Macie administrator account to enable Macie. Configure automated sensitive data discovery for all accounts in the organization.

C. For each Region in the Organizations management account, configure a service control policy (SCP) to identify PII. Apply the SCP to the organization root.

D. In the Organizations management account, configure AWS Lambda functions to scan for PII in each Region.

**Answer: B**

**Explanation:** Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data in AWS. To scale across Regions and accounts in AWS Organizations, Macie supports delegated administration, automated sensitive data discovery, and multi-account aggregation through a centralized admin account.

**QUESTION 1258** A company is planning to deploy a data processing platform on AWS. The data processing platform is based on PostgreSQL. The company stores the data that the platform must process on-premises. To comply with regulations, the company must not migrate the data to the cloud. However, the company wants to use AWS managed data analytics solutions. Which solution will meet these requirements?

A. Create an Amazon RDS for PostgreSQL database in a VPC. Create an interface VPC endpoint to connect the on-premises PostgreSQL database to the RDS for PostgreSQL database.

B. Create Amazon EC2 instances in an Auto Scaling group on AWS Outposts. Install PostgreSQL data analytics software on the instances.

C. Create an Amazon EMR cluster on AWS Outposts. Connect the EMR cluster to the on-premises PostgreSQL database to perform data processing locally.

D. Create an Amazon EMR cluster in a VPC. Connect the EMR cluster to Amazon RDS for SQL Server with a linked server to connect to the company's data processing platform.

**Answer: C**

**Explanation:** AWS Outposts extends AWS infrastructure and services to on-premises locations. Running Amazon EMR on Outposts allows for processing data that resides locally while benefiting from the managed services of EMR. This enables compliance with data residency requirements and provides scalability and manageability for analytics.

**QUESTION 1259** A company runs a critical three-tier web application that consists of multiple virtual machines (VMs) and virtual databases in an on-premises environment. The company wants to set up a disaster recovery (DR) environment in AWS. The company requires a 15-minute recovery time objective (RTO). The company must be able to test the failover solution to validate the recovery. The solution must provide an automated failover mechanism. Which solution will meet these requirements?

A. Use AWS Backup to create backups of the on-premises VMs and to restore the backups in AWS. Configure recovery to Amazon EC2 instances to meet the RTO requirement.

B. Use AWS Database Migration Service (AWS DMS) to replicate the on-premises databases to Amazon RDS. Set up AWS Storage Gateway for baseline and incremental data replication to AWS to meet the RTO requirement.

C. Use AWS DataSync and AWS Storage Gateway to migrate the baseline and incremental data to AWS. Use Amazon EC2, Amazon S3, and an Application Load Balancer to set up the DR environment.

D. Use AWS Elastic Disaster Recovery to replicate the VMs incrementally to AWS. Configure Elastic Disaster Recovery to automate the DR process.

**Answer: D**

**Explanation:** AWS Elastic Disaster Recovery (AWS DRS) enables fast, reliable, and cost-effective disaster recovery. It replicates on-premises machines to AWS using continuous block-level replication. It supports automated testing and failover, meeting aggressive RTO targets such as 15 minutes.

**QUESTION 1260** A company collects data from sensors. The company needs a cloud-based solution to store and transform the sensor data to make critical decisions. The solution must store the data for up to 2 days. After 2 days, the solution must delete the data. The company needs to use the transformed data in an automated workflow that has manual approval steps. Which solution will meet these requirements?

A. Load the data into an Amazon Simple Queue Service (Amazon SQS) queue that has a retention period of 2 days. Use an Amazon

EventBridge pipe to retrieve data from the queue, transform the data, and pass the data to an AWS Step Functions workflow.B. Load the data into AWS DataSync. Delete the DataSync task after 2 days. Invoke an AWS Lambda function to retrieve the data, transform the data, and invoke a second Lambda function that performs the remaining workflow steps.C. Load the data into an Amazon Simple Notification Service (Amazon SNS) topic. Use an Amazon EventBridge pipe to retrieve the data from the topic, transform the data, and send the data to Amazon EC2 instances to perform the remaining workflow steps.D. Load the data into an Amazon Simple Notification Service (Amazon SNS) topic. Use an Amazon EventBridge pipe to retrieve the data from the topic and transform the data into an appropriate format for an Amazon SQS queue. Use an AWS Lambda function to poll the queue to perform the remaining workflow steps.

**Answer:** AExplanation:Amazon SQS with a 2-day retention ensures the data lives just as long as needed. EventBridge Pipes allow direct integration between event producers and consumers, with optional filtering and transformation. AWS Step Functions supports manual approval steps, which fits the workflow requirement perfectly.

**QUESTION 1261** A company hosts an application on AWS that gives users the ability to download photos. The company stores all photos in an Amazon S3 bucket that is located in the us-east-1 Region. The company wants to provide the photo download application to global customers with low latency. Which solution will meet these requirements?

A. Find the public IP addresses that Amazon S3 uses in us-east-1. Configure an Amazon Route 53 latency-based routing policy that routes to all the public IP addresses.

B. Configure an Amazon CloudFront distribution in front of the S3 bucket. Use the distribution endpoint to access the photos that are in the S3 bucket.

C. Configure an Amazon Route 53 geoproximity routing policy to route the traffic to the S3 bucket that is closest to each customer's location.

D. Create a new S3 bucket in the us-west-1 Region. Configure an S3 Cross-Region Replication rule to copy the photos to the new S3 bucket.

**Answer:** BExplanation:Amazon CloudFront is a content delivery network (CDN) service that distributes content with low latency and high transfer speeds. Placing CloudFront in front of the S3 bucket ensures global users download content from the nearest edge location, reducing latency significantly.

**QUESTION 1262** A company has a three-tier environment on AWS that ingests sensor data from its users' devices. The traffic flows through a Network Load Balancer (NLB) then to Amazon EC2 instances for the web tier and finally to EC2 instances for the application tier that makes database calls. What should a solutions architect do to improve the security of data in transit to the web tier?

A. Configure a TLS listener and add the server certificate on the NLB.

B. Configure AWS Shield Advanced and enable AWS WAF on the NLB.

C. Change the load balancer to an Application Load Balancer and attach AWS WAF to it.

D. Encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances using AWS Key Management Service (AWS KMS).

**Answer:** AExplanation:Implement secure key and certificate management: Store encryption keys and certificates securely and rotate them at appropriate time intervals while applying strict access control; for example, by using a certificate management service, such as AWS Certificate Manager (ACM).Enforce encryption in transit: Enforce your defined encryption requirements based on appropriate standards and recommendations to help you meet your organizational, legal, and compliance requirements.

Automate detection of unintended data access: Use tools such as GuardDuty to automatically detect attempts to move data outside of defined boundaries based on data classification level, for example, to detect a trojan that is copying data to an unknown or untrusted network using the DNS protocol.

Authenticate network communications: Verify the identity of communications by using protocols that support authentication, such as Transport Layer Security (TLS) or IPsec.

**QUESTION 1263** A company is building a serverless web application with multiple interdependent workflows that millions of users worldwide will access. The application needs to handle bursts of traffic. Which solution will meet these requirements MOST cost-effectively?

A. Deploy an Amazon API Gateway HTTP API with a usage plan and throttle settings.

B. Deploy an Amazon API Gateway HTTP API with a usage plan and throttle settings. Use AWS Step Functions with a Standard Workflow.

C. Deploy an Amazon API Gateway HTTP API without a usage plan. Use AWS Step Functions with an Express Workflow.

D. Deploy an Amazon API Gateway HTTP API without a usage plan. Use AWS Step Functions and multiple AWS Lambda functions with reserved concurrency.

**Answer:** BExplanation:Express Workflows in AWS Step Functions are optimized for high-throughput, short-duration, and low-cost workflows. They are suitable for applications with large volumes of parallel and interdependent tasks. When paired with HTTP APIs from API Gateway, which are more cost-efficient than REST APIs, this setup offers scalability and cost-effectiveness.

**QUESTION 1264** A company is hosting multiple websites for several lines of business under its registered parent domain. Users accessing these websites will be routed to appropriate backend Amazon EC2 instances based on the subdomain. The websites host static webpages, images, and server-side scripts like PHP and JavaScript. Some of the websites experience peak access during the first two hours of business with constant usage throughout the rest of the day. A solutions architect needs to design a solution that will automatically adjust capacity to these traffic patterns while keeping costs low. Which combination of AWS services or features will meet these requirements? (Choose two.)

A. AWS Batch

B. Network Load Balancer

C. Application Load Balancer

D. Amazon EC2 Auto Scaling

E. Amazon S3 website hosting

**Answer:** CDExplanation:An Application Load Balancer supports path- and host-based routing, which makes it ideal

for routing requests based on subdomains. EC2 Auto Scaling ensures that the number of instances adjusts dynamically based on traffic, which helps manage cost and performance during predictable peak hours.

**QUESTION 1265** An internal product team is deploying a new application to a private VPC in a company's AWS account. The application runs on Amazon EC2 instances that are in a security group named App1. The EC2 instances store application data in an Amazon S3 bucket and use AWS Secrets Manager to store application service credentials. The company's security policy prohibits applications in a private VPC from using public IP addresses to communicate. Which combination of solutions will meet these requirements? (Choose two.)

A. Configure gateway endpoints for Amazon S3 and AWS Secrets Manager.

B. Configure interface VPC endpoints for Amazon S3 and AWS Secrets Manager.

C. Add routes to the endpoints in the VPC route table.

D. Associate the App1 security group with the interface VPC endpoints. Configure a self-referencing security group rule to allow inbound traffic.

E. Associate the App1 security group with the gateway endpoints. Configure a self-referencing security group rule to allow inbound traffic.

**Answer: BCE**

**Explanation:** To securely access AWS services like S3 and Secrets Manager from a private VPC without using public IPs, interface VPC endpoints are required. These endpoints are accessible via private IP addresses. For the application to reach these endpoints, appropriate routes must be configured in the route table.

**QUESTION 1266** A company wants to optimize costs for its AWS infrastructure. The company wants to receive notifications when actual costs or forecasted costs exceed a specified budget. The company does not want to develop a custom solution. Which solution will meet these requirements?

A. Use AWS Trusted Advisor to set up budget notifications. Configure Amazon CloudWatch to monitor costs. Export CloudWatch data to Amazon S3. Use machine learning (ML) to estimate future trends based on the CloudWatch data.

B. Create a budget in AWS Budgets that has a specified cost threshold. Create an AWS Lambda function that sends a notification to the company when costs reach the specified threshold. Use AWS Billing and Cost Management reports to monitor costs.

C. Use AWS Cost Explorer to set a specified budget threshold. Create an AWS Lambda function to calculate cost estimates. Configure the Lambda function to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic if estimated costs exceed the specified threshold.

D. Create a budget in AWS Budgets that has a specified cost threshold. Configure AWS Budgets to send budget alerts to an Amazon Simple Notification Service (Amazon SNS) topic. Use AWS Cost Explorer to monitor costs.

**Answer: D**

**Explanation:** AWS Budgets allows you to set custom cost and usage budgets. When actual or forecasted usage exceeds the threshold, AWS Budgets can automatically send alerts to an Amazon SNS topic. You can use AWS Cost Explorer in parallel for visual tracking of spending. This solution requires no code and has minimal operational overhead.

**QUESTION 1267** A company launches a new web application that uses an Amazon Aurora PostgreSQL database. The company wants to add new features to the application that rely on AI. The company requires vector storage capability to use AI tools. Which solution will meet this requirement MOST cost-effectively?

A. Use Amazon OpenSearch Service to create an OpenSearch service. Configure the application to write vector embeddings to a vector index.

B. Create an Amazon DocumentDB cluster. Configure the application to write vector embeddings to a vector index.

C. Create an Amazon Neptune ML cluster. Configure the application to write vector embeddings to a vector graph.

D. Install the pgvector extension on the Aurora PostgreSQL database. Configure the application to write vector embeddings to a vector table.

**Answer: D**

**Explanation:** Aurora PostgreSQL supports the pgvector extension, which allows storage and querying of vector embeddings directly inside the database. This eliminates the need for external vector databases and provides cost-effective and performant integration for AI workloads.

**QUESTION 1268** A company manages multiple AWS accounts in an organization in AWS Organizations. The company's applications run on Amazon EC2 instances in multiple AWS Regions. The company needs a solution to simplify the management of security rules across the accounts in its organization. The solution must apply shared security group rules, audit security groups, and detect unused and redundant rules in VPC security groups across all AWS environments. Which solution will meet these requirements with the MOST operational efficiency?

A. Use AWS Firewall Manager to create a set of rules based on the security requirements. Replicate the rules to all the AWS accounts and Regions.

B. Use AWS CloudFormation StackSets to provision VPC security groups based on the specifications across multiple accounts and Regions. Deploy AWS Network Firewall to define the firewall rules to control network traffic across multiple accounts and Regions.

C. Use AWS CloudFormation StackSets to provision VPC security groups based on the specifications across multiple accounts and Regions. Configure AWS Config and AWS Lambda to evaluate compliance information and to automate enforcement across all accounts and Regions.

D. Use AWS Network Firewall to build policies based on the security requirements. Centrally apply the new policies to all the VPCs and accounts.

**Answer: A**

**Explanation:** AWS Firewall Manager integrates with AWS Organizations to centrally manage and apply security group policies, AWS WAF rules, and AWS Shield Advanced protections. It automates the propagation of rules across accounts and Regions and can also audit and remediate noncompliant configurations.

**QUESTION 1269** A company uses Amazon EC2 instances and Amazon Elastic Block Store (Amazon EBS) volumes to run an application. The company creates one snapshot of each EBS volume every day. The company needs to prevent users from accidentally deleting the EBS volume snapshots. The solution must not change the

administrative rights of a storage administrator user. Which solution will meet these requirements with the LEAST administrative effort?A. Create an IAM role that has permission to delete snapshots. Attach the role to a new EC2 instance. Use the AWS CLI from the new EC2 instance to delete snapshots.B. Create an IAM policy that denies snapshot deletion. Attach the policy to the storage administrator user.C. Add tags to the snapshots. Create tag-level retention rules in the Recycle Bin for EBS snapshots. Configure rule lock settings for the retention rules.D. Take EBS snapshots by using the EBS direct APIs. Copy the snapshots to an Amazon S3 bucket. Configure S3 Versioning and Object Lock on the bucket. Answer: CExplanation:Amazon EBS Snapshots Recycle Bin enables you to specify retention rules for EBS snapshots based on tags. When snapshots are deleted, they are retained in the Recycle Bin for a specified duration, preventing accidental deletion. Tag-level rules allow selective protection without changing IAM roles or user permissions. QUESTION 1270A company is designing an application to connect AWS Lambda functions to an Amazon RDS for MySQL DB instance. The DB instance manages many connections. The company needs to modify the application to improve connectivity and recovery. Which solution will meet these requirements with the LEAST operational overhead?A. Use Amazon RDS Proxy for connection pooling. Modify the application to use the RDS Proxy for connections to the DB instance.B. Create a new RDS instance for connection pooling. Modify the application to use the new RDS instance for connectivity.C. Create read replicas to distribute the load of the DB instance. Create a Network Load Balancer to distribute the load across the read replicas.D. Migrate the RDS for MySQL DB instance to Amazon Aurora MySQL to increase DB instance performance. Answer: AExplanation:Amazon RDS Proxy helps manage thousands of concurrent database connections by pooling and reusing them efficiently. It is especially useful for serverless applications like AWS Lambda that can open numerous connections quickly, potentially overwhelming the database. Using RDS Proxy reduces connection management overhead and improves fault tolerance. QUESTION 1271A company has an application that processes information from documents that users upload. When a user uploads a new document to an Amazon S3 bucket, an AWS Lambda function is invoked. The Lambda function processes information from the documents. The company discovers that the application did not process many recently uploaded documents. The company wants to ensure that the application processes each document with retries if there is an error during the first attempt to process the document. Which solution will meet these requirements?A. Create an Amazon API Gateway REST API that has a proxy integration to the Lambda function. Update the application to send requests to the REST API.B. Configure a replication policy on the S3 bucket to stage the documents in another S3 bucket that an AWS Batch job processes on a daily schedule.C. Deploy an Application Load Balancer in front of the Lambda function that processes the documents.D. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an event source for the Lambda function. Configure an S3 event notification on the S3 bucket to send new document upload events to the SQS queue. Answer: DExplanation:Using SQS as a buffer between S3 and the Lambda function ensures durability and allows for retries in case of processing failures. Messages in the queue can be retried by Lambda, and failed processing can be directed to a dead-letter queue for further inspection. This guarantees reliable and scalable message-driven processing. QUESTION 1272A company's expense tracking application gives users the ability to upload images of receipts. The application analyzes the receipts to extract information and stores the raw images in Amazon S3. The application is written in Java and runs on Amazon EC2 On-Demand Instances in an Auto Scaling group behind an Application Load Balancer. The compute costs and storage costs have increased with the popularity of the application. Which solution will provide the MOST cost savings without affecting application performance?A. Purchase a Compute Savings Plan for the maximum number of necessary EC2 instances. Store the uploaded files in Amazon Elastic File System (Amazon EFS).B. Decrease the minimum number of EC2 instances in the Auto Scaling group. Use On-Demand Instances for peak scaling. Store the uploaded files in Amazon Elastic File System (Amazon EFS).C. Decrease the maximum number of EC2 instances in the Auto Scaling group. Set up S3 Lifecycle policies to archive the raw images to lower-cost storage tiers after 30 days.D. Purchase a Compute Savings Plan for the minimum number of necessary EC2 instances. Use On-Demand Instances for peak scaling. Set up S3 Lifecycle policies to archive the raw images to lower-cost storage tiers after 30 days. Answer: DExplanation:Purchasing a Compute Savings Plan for the minimum baseline usage ensures cost savings. Using On-Demand Instances for peak times ensures flexibility without over-provisioning. S3 Lifecycle policies enable automatic transition of objects to lower-cost storage classes such as S3 Glacier or S3 Intelligent-Tiering, further reducing storage costs. QUESTION 1273A solutions architect is designing the architecture for a web application that has a frontend and a backend. The backend services must receive data from the frontend services for processing. The frontend must manage access to the application by using API keys. The backend must scale without affecting the frontend. Which solution will meet these requirements?A. Deploy an Amazon API Gateway HTTP API as the frontend to direct traffic to an Amazon Simple Queue Service (Amazon SQS) queue. Use AWS Lambda functions as the backend to read from the queue.B. Deploy an Amazon API Gateway REST API as the frontend to direct traffic to an Amazon Simple Queue Service (Amazon SQS) queue. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate as the backend to read from the queue.C. Deploy an Amazon API

Gateway REST API as the frontend to direct traffic to an Amazon Simple Notification Service (Amazon SNS) topic. Use AWS Lambda functions as the backend. Subscribe the Lambda functions to the topic.

D. Deploy an Amazon API Gateway HTTP API as the frontend to direct traffic to an Amazon Simple Notification Service (Amazon SNS) topic. Use Amazon Elastic Kubernetes Service (Amazon EKS) on AWS Fargate as the backend. Subscribe Amazon EKS to the topic.

Answer: AExplanation: Using API Gateway with API keys provides secure access control. Amazon SQS allows asynchronous decoupling between frontend and backend, ensuring that backend processing can scale independently. AWS Lambda reading from SQS ensures scalable, event-driven processing with minimal operational management. This architecture is resilient and decoupled.

QUESTION 1274 A company is migrating a daily Microsoft Windows batch job from the company's on-premises environment to AWS. The current batch job runs for up to 1 hour. The company wants to modernize the batch job process for the cloud environment. Which solution will meet these requirements with the LEAST operational overhead?

A. Create a fleet of Amazon EC2 instances in an Auto Scaling group to handle the Windows batch job processing.

B. Implement an AWS Lambda function to process the Windows batch job. Use an Amazon EventBridge rule to invoke the Lambda function.

C. Use AWS Fargate to deploy the Windows batch job as a container. Use AWS Batch to manage the batch job processing.

D. Use Amazon Elastic Kubernetes Service (Amazon EKS) on Amazon EC2 instances to orchestrate Windows containers for the batch job processing.

Answer: CExplanation: AWS Batch supports Windows-based jobs and automates provisioning and scaling of compute environments. Paired with AWS Fargate, it removes the need to manage infrastructure. This solution requires the least operational overhead and is cloud-native, providing flexibility and scalability.

QUESTION 1275 A company is building a serverless application that processes large volumes of data from a mobile app. The application uses an AWS Lambda function to process the data and store the data in an Amazon DynamoDB table. The company needs to ensure that the application can recover from failures and continue processing data without losing any records. Which solution will meet these requirements?

A. Configure the Lambda function to use a dead-letter queue with an Amazon Simple Queue Service (Amazon SQS) queue. Configure Lambda to retry failed records from the dead-letter queue. Use a retry mechanism by implementing an exponential backoff algorithm.

B. Configure the Lambda function to read records from Amazon Data Firehose. Replay the Firehose records in case of any failures.

C. Use Amazon OpenSearch Service to store failed records. Configure AWS Lambda to retry failed records from OpenSearch Service. Use Amazon EventBridge to orchestrate the retry logic.

D. Use Amazon Simple Notification Service (Amazon SNS) to store the failed records. Configure Lambda to retry failed records from the SNS topic. Use Amazon API Gateway to orchestrate the retry calls.

Answer: AExplanation: Dead-letter queues (DLQs) with Amazon SQS allow Lambda functions to offload failed events for later inspection or retry. Using retry logic with exponential backoff ensures resilience and compliance with best practices for fault-tolerant serverless architectures. This guarantees no data is lost due to transient errors.

QUESTION 1276 A company discovers that an Amazon DynamoDB Accelerator (DAX) cluster for the company's web application workload is not encrypting data at rest. The company needs to resolve the security issue. Which solution will meet this requirement?

A. Stop the existing DAX cluster. Enable encryption at rest for the existing DAX cluster, and start the cluster again.

B. Delete the existing DAX cluster. Recreate the DAX cluster, and configure the new cluster to encrypt the data at rest.

C. Update the configuration of the existing DAX cluster to encrypt the data at rest.

D. Integrate the existing DAX cluster with AWS Security Hub to automatically enable encryption at rest.

Answer: BExplanation: DAX does not support enabling encryption at rest on an existing cluster. To use encryption at rest, you must create a new DAX cluster with encryption enabled at creation time and migrate workloads accordingly.

QUESTION 1277 A company is developing a serverless web application that gives users the ability to interact with real-time analytics from online games. The data from the games must be streamed in real time. The company needs a durable, low-latency database option for user data. The company does not know how many users will use the application. Any design considerations must provide response times of single-digit milliseconds as the application scales. Which combination of AWS services will meet these requirements? (Choose two.)

A. Amazon CloudFront

B. Amazon DynamoDB

C. Amazon Kinesis

D. Amazon RDSE

AWS Global Accelerator

Answer: BCExplanation: Amazon Kinesis allows real-time ingestion of game events at scale, while Amazon DynamoDB provides millisecond-latency access to user data, automatically scaling with demand. This combination ensures real-time processing and fast data retrieval without managing infrastructure.

QUESTION 1278 A company is migrating a production environment application to the AWS Cloud. The company uses Amazon RDS for Oracle for the database layer. The company needs to configure the database to meet the needs of high I/O intensive workloads that require low latency and consistent throughput. The database workloads are read intensive and write intensive. Which solution will meet these requirements?

A. Use a Multi-AZ DB instance deployment for the RDS for Oracle database.

B. Configure the RDS for Oracle database to use the Provisioned IOPS SSD storage type.

C. Configure the RDS for Oracle database to use the General Purpose SSD storage type.

D. Enable RDS read replicas for RDS for Oracle.

Answer: BExplanation: Provisioned IOPS SSD (io1 or io2) is designed for I/O-intensive workloads that require low latency and consistent throughput, which is critical for transactional and

production databases. It provides predictable performance, unlike General Purpose SSD, which is burst-based. **QUESTION 1279** A company hosts a public web application on AWS. The website has a three-tier architecture. The frontend web tier is comprised of Amazon EC2 instances in an Auto Scaling group. The application tier is a second Auto Scaling group. The database tier is an Amazon RDS database. The company has configured the Auto Scaling groups to handle the application's normal level of demand. During an unexpected spike in demand, the company notices a long delay in the startup time when the frontend and application layers scale out. The company needs to improve the scaling performance of the application without negatively affecting the user experience. Which solution will meet these requirements **MOST** cost-effectively?  
A. Decrease the minimum number of EC2 instances for both Auto Scaling groups. Increase the desired number of instances to meet the peak demand requirement.  
B. Configure the maximum number of instances for both Auto Scaling groups to be the number required to meet the peak demand.  
C. Create a warm pool.  
D. Increase the maximum number of EC2 instances for both Auto Scaling groups to meet the normal demand requirement. Create a warm pool.  
**Answer: B**  
**Explanation:** EC2 Auto Scaling warm pools allow you to pre-initialize instances, reducing the delay in scale-out events. This results in significantly faster response times during demand surges while remaining cost-effective compared to always running at peak capacity.

**QUESTION 1280** A telemarketing company is designing its customer call center functionality on AWS. The company needs a solution that provides multiple speaker recognition and generates transcript files. The company wants to query the transcript files to analyze the business patterns. Which solution will meet these requirements?  
A. Use Amazon Rekognition for multiple speaker recognition. Store the transcript files in Amazon S3. Use machine learning (ML) models to analyze the transcript files.  
B. Use Amazon Transcribe for multiple speaker recognition. Use Amazon Athena to analyze the transcript files.  
C. Use Amazon Translate for multiple speaker recognition. Store the transcript files in Amazon Redshift. Use SQL queries to analyze the transcript files.  
D. Use Amazon Rekognition for multiple speaker recognition. Store the transcript files in Amazon S3. Use Amazon Textract to analyze the transcript files.  
**Answer: B**  
**Explanation:** Amazon Transcribe supports automatic speech recognition (ASR) with speaker diarization (i.e., multiple speaker identification). The transcripts can be stored in Amazon S3 and queried using Amazon Athena, which provides a serverless, pay-as-you-go interactive querying model.

**QUESTION 1281** A company has a three-tier web application. An Application Load Balancer (ALB) is in front of Amazon EC2 instances that are in the ALB target group. An Amazon S3 bucket stores documents. The company requires the application to meet a recovery time objective (RTO) of 60 seconds. Which solution will meet this requirement?  
A. Replicate S3 objects to a second AWS Region. Create a second ALB and a minimum set of EC2 instances in the second Region. Ensure that the EC2 instances are shut down until they are needed. Configure Amazon Route 53 to fail over to the second Region by using an IP-based routing policy.  
B. Use AWS Backup to take hourly backups of the EC2 instances. Back up the S3 data to a second AWS Region. Use AWS CloudFormation to deploy the entire infrastructure in the second Region when needed.  
C. Create daily snapshots of the EC2 instances in a second AWS Region. Use the snapshots to recreate the instances in the second Region. Back up the S3 data to the second Region. Perform a failover by modifying the application DNS record when needed.  
D. Replicate S3 objects to a second AWS Region. Create a second ALB and a minimum set of EC2 instances in the second Region. Ensure that the EC2 instances in the second Region are running. Configure Amazon Route 53 to fail over to the secondary Region based on health checks.  
**Answer: D**  
**Explanation:** To achieve a 60-second RTO, pre-warming the DR environment (including running EC2 instances and Route 53 health checks) is essential. Active/passive failover using Route 53 with health checks ensures fast redirection when the primary Region becomes unavailable. S3 cross-region replication ensures document availability.

**QUESTION 1282** A media company runs an application on multiple Amazon EC2 instances that requires high storage input/output operations per second (IOPS). To achieve the necessary performance, a solutions architect wants to stripe multiple Amazon EBS volumes together and attach the volumes to EC2 instances. The solutions architect wants to receive a notification when IOPS are over-provisioned. Which solution will meet these requirements?  
A. Configure auto scaling for the EBS volumes to automatically increase or decrease IOPS based on the EC2 instance CPU utilization metric.  
B. Deploy the application on an EC2 instance type that supports the highest possible IOPS.  
C. Create a custom AWS Config rule to monitor the provisioned IOPS for the EBS volumes that are attached to the EC2 instances and to send notifications.  
D. Adjust the IOPS of each EBS volume daily based on Amazon CloudWatch metrics for IOPS utilization.  
**Answer: C**  
**Explanation:** AWS Config allows for creation of custom rules to monitor EBS configurations. Combined with CloudWatch metrics and Amazon SNS, custom rules can track over-provisioned IOPS and send alerts when thresholds are breached, allowing proactive cost and performance management.

**QUESTION 1283** An insurance company runs an application on premises to process contracts. The application processes jobs that are comprised of many tasks. The individual tasks run for up to 5 minutes. Some jobs can take up to 24 hours in total to finish. If a task fails, the task must be reprocessed. The company wants to migrate the application to AWS. The company will use Amazon S3 as part of the solution. The company wants to configure jobs to start

automatically when a contract is uploaded to an S3 bucket. Which solution will meet these requirements? A. Use AWS Lambda functions to process individual tasks. Create a primary Lambda function to handle the overall job processing by calling individual Lambda functions in sequence. Configure the S3 bucket to send an event notification to invoke the primary Lambda function to begin processing. B. Use a state machine in AWS Step Functions to handle the overall contract processing job. Configure the S3 bucket to send an event notification to Amazon EventBridge. Create a rule in Amazon EventBridge to target the state machine. C. Use an AWS Batch job to handle the overall contract processing job. Configure the S3 bucket to send an event notification to initiate the Batch job. D. Use an S3 event notification to notify an Amazon Simple Queue Service (Amazon SQS) queue when a contract is uploaded. Configure an AWS Lambda function to read messages from the queue and to run the contract processing job.

**Answer: B**

**Explanation:** AWS Step Functions supports long-running workflows and error retries, making it ideal for a job composed of many tasks. Integration with EventBridge allows automatic triggering from S3 events. This setup is resilient and supports up to 1-year execution duration.

**QUESTION 1284** A company wants to visualize its AWS spend and resource usage. The company wants to use an AWS managed service to provide visual dashboards. Which solution will meet these requirements? A. Configure an export in AWS Data Exports. Use Amazon QuickSight to create a cost and usage dashboard. View the data in QuickSight. B. Configure one custom budget in AWS Budgets for costs. Configure a second custom budget for usage. Schedule daily AWS Budgets reports by using the two budgets as sources. C. Configure AWS Cost Explorer to use user-defined cost allocation tags with hourly granularity to generate detailed data. D. Configure an export in AWS Data Exports. Use the standard export option. View the data in Amazon Athena.

**Answer: A**

**Explanation:** By exporting AWS Cost and Usage Reports (CUR) to Amazon S3 and analyzing them with Amazon QuickSight, companies can generate interactive visual dashboards. This solution is fully AWS-managed, requires no third-party tools, and integrates deeply with AWS cost data.

**QUESTION 1285** A company hosts an application on AWS. The application has generated approximately 2.5 TB of data over the previous 12 years. The company currently stores the data on Amazon EBS volumes. The company wants a cost-effective backup solution for long-term storage. The company must be able to retrieve the data within minutes when required for audits. Which solution will meet these requirements? A. Create EBS snapshots to back up the data. B. Create an Amazon S3 bucket. Use the S3 Glacier Deep Archive storage class to back up the data. C. Create an Amazon S3 bucket. Use the S3 Glacier Flexible Retrieval storage class to back up the data. D. Create an Amazon Elastic File System (Amazon EFS) file system to back up the data.

**Answer: C**

**Explanation:** Amazon S3 Glacier Flexible Retrieval is a low-cost archival storage class that supports retrieval of data within minutes (expedited access ~1-5 minutes), making it ideal for audit scenarios where occasional, quick access to archived data is required. In contrast, Glacier Deep Archive takes hours to retrieve.

**QUESTION 1286** A company is developing a new online gaming application. The application will run on Amazon EC2 instances in multiple AWS Regions and will have a high number of globally distributed users. A solutions architect must design the application to optimize network latency for the users. Which actions should the solutions architect take to meet these requirements? (Choose two.)

A. Configure AWS Global Accelerator. Create Regional endpoint groups in each Region where an EC2 fleet is hosted. B. Create a content delivery network (CDN) by using Amazon CloudFront. Enable caching for static and dynamic content, and specify a high expiration period. C. Integrate AWS Client VPN into the application. Instruct users to select which Region is closest to them after they launch the application. Establish a VPN connection to that Region. D. Create an Amazon Route 53 weighted routing policy. Configure the routing policy to give the highest weight to the EC2 instances in the Region that has the largest number of users.

E. Configure an Amazon API Gateway endpoint in each Region where an EC2 fleet is hosted. Instruct users to select which Region is closest to them after they launch the application. Use the API Gateway endpoint that is closest to them.

**Answer: AB**

**Explanation:** AWS Global Accelerator reduces latency by directing users to the optimal Regional endpoint based on global network health and proximity. Amazon CloudFront caches static and dynamic content at edge locations for ultra-low latency access worldwide, improving performance and reducing server load.

**QUESTION 1287** A company runs a web application in a single AWS Region. A solutions architect wants to ensure that the web application can continue to operate if the application becomes unavailable in the Region. Which solution will meet this requirement? A. Deploy the application in multiple Regions. Use Amazon Route 53 DNS health checks to route traffic to a healthy Region. B. Deploy the application in multiple Availability Zones within a single Region. Use Amazon Route 53 DNS health checks to route traffic to healthy application resources. C. Deploy the application in multiple Regions. Use an Amazon Route 53 simple routing record to route traffic to a healthy Region. D. Deploy the application in multiple Availability Zones within a single Region. Use an Amazon Route 53 latency record in each Availability Zone to route traffic to a healthy Availability Zone.

**Answer: A**

**Explanation:** To protect against a Regional failure, the application must be deployed in multiple Regions. Amazon Route 53 DNS failover with health checks allows traffic to be automatically routed to a healthy Region when the primary becomes unavailable, meeting high availability and disaster recovery requirements.

**QUESTION 1288** An application uses an Amazon SQS queue and two AWS Lambda functions. One of the Lambda functions pushes messages to the queue, and the other

function polls the queue and receives queued messages. A solutions architect needs to ensure that only the two Lambda functions can write to or read from the queue. Which solution will meet these requirements?

A. Attach an IAM policy to the SQS queue that grants the Lambda function principals read and write access. Attach an IAM policy to the execution role of each Lambda function that denies all access to the SQS queue except for the principal of each function.

B. Attach a resource-based policy to the SQS queue to deny read and write access to the queue for any entity except the principal of each Lambda function. Attach an IAM policy to the execution role of each Lambda function that allows read and write access to the queue.

C. Attach a resource-based policy to the SQS queue that grants the Lambda function principals read and write access to the queue. Attach an IAM policy to the execution role of each Lambda function that allows read and write access to the queue.

D. Attach a resource-based policy to the SQS queue to deny all access to the queue. Attach an IAM policy to the execution role of each Lambda function that grants read and write access to the queue.

Answer: C

Explanation: To ensure that only specific AWS Lambda functions can read from or write to an Amazon SQS queue, user resource-based policies are attached directly to the SQS queue. These policies explicitly grant permissions to the IAM roles used by the Lambda functions. Additionally, the Lambda execution roles must also have IAM policies that permit SQS access. This dual-layer approach follows the AWS security best practice of granting least privilege access and ensures that no other service or entity can interact with the queue. This is a common and supported pattern documented in the Amazon SQS Developer Guide, where resource-based policies restrict access at the queue level while IAM roles control permissions at the function level.

**QUESTION 1289**

A company wants to migrate its on-premises Oracle database to Amazon Aurora. The company wants to use a secure and encrypted network to transfer the data. Which combination of steps will meet these requirements? (Choose two.)

A. Use AWS Application Migration Service to migrate the data.

B. Use AWS Schema Conversion Tool (AWS SCT) and AWS Database Migration Service (AWS DMS) to migrate the data.

C. Use AWS Direct Connect SiteLink to transfer data from the on-premises environment to AWS.

D. Use AWS Site-to-Site VPN to establish a connection to transfer the data from the on-premises environment to AWS.

E. Use AWS App2Container to migrate the data.

Answer: BD

**QUESTION 1290**

A company has a single AWS account that contains resources belonging to several teams. The company needs to identify the costs associated with each team. The company wants to use a tag named **CostCenter** to identify resources that belong to each team. Which of the following actions should the company take to meet this requirement? (Choose two.)

A. Tag all resources that belong to each team with the user-defined **CostCenter** tag.

B. Create a tag for each team, and set the value to **CostCenter**.

C. Activate the **CostCenter** tag to track cost allocation.

D. Configure AWS Billing and Cost Management to send monthly invoices to the company through email messages.

E. Set up consolidated billing in the existing AWS account.

Answer: AC

**QUESTION 1291**

A financial services company must retain log data for 1 year. The company stores log files in an Amazon S3 bucket and wants to prevent any user from deleting or overwriting the log files during this period. The data must remain available for read-only requests. Which solution will meet these requirements?

A. Enable S3 Versioning on the bucket. Use Object Lock in compliance mode with a 1-year retention period.

B. Enable S3 Transfer Acceleration on the bucket. Create an S3 Lifecycle Configuration rule to move objects to Amazon S3 Glacier Flexible Retrieval after 1 year.

C. Enable S3 Versioning on the bucket. Create an S3 Lifecycle Configuration rule to move objects to Amazon S3 Glacier Flexible Retrieval after 1 year.

D. Create an AWS Lambda function to programmatically check the timestamp of S3 data and to move the data to Amazon S3 Glacier Deep Archive if the data is older than 1 year.

Answer: A

**QUESTION 1292**

A company has an industrial application that controls a process in real time. The company plans to rearchitect the application to distribute jobs across several Amazon EC2 instances in a VPC. The solution needs to maximize the network throughput and minimize the network latency between the instances. Which solution will meet these requirements?

A. Place the instances in a host-level partition placement group. Choose instance types that support enhanced networking.

B. Place the instances in several dedicated hosts in the same partition of a partition placement group. Choose dedicated hosts that support enhanced networking.

C. Place the instances in several dedicated hosts in the same rack of a rack-level placement group. Choose dedicated hosts that support enhanced networking.

D. Place the instances in a cluster placement group. Choose instance types that support enhanced networking.

Answer: D

**QUESTION 1293**

A company is planning to migrate customer records to an Amazon S3 bucket. The company needs to ensure that customer records are protected against unauthorized access and are encrypted in transit and at rest. The company must monitor all access to the S3 bucket. Which solution will meet these requirements?

A. Use AWS Key Management Service (AWS KMS) to encrypt customer records at rest. Create an S3 bucket policy that includes the `aws:SecureTransport` condition. Use an IAM policy to control access to the records. Use AWS CloudTrail to monitor access to the records.

B. Use AWS Nitro Enclaves to encrypt customer records at rest. Use AWS Key Management Service (AWS KMS) to encrypt the records in transit. Use an IAM policy to control access to the records. Use AWS CloudTrail and AWS Security Hub to monitor access to the records.

C. Use AWS Key Management Service (AWS KMS) to encrypt customer records at rest. Create an Amazon Cognito user pool to control access to the records. Use AWS CloudTrail to monitor access to the records. Use Amazon GuardDuty to detect threats.

D. Use server-side

encryption with Amazon S3 managed keys (SSE-S3) with default settings to encrypt the records at rest. Access the records by using an Amazon CloudFront distribution that uses the S3 bucket as the origin. Use IAM roles to control access to the records. Use Amazon CloudWatch to monitor access to the records.

Answer: AQUESTION 1294A company plans to deploy containerized microservices in the AWS Cloud. The containers must mount a persistent file store that the company can manage by using OS-level permissions. The company requires fully managed services to host the containers and file store. Which solution will meet these requirements?

A. Use AWS Lambda functions and an Amazon API Gateway REST API to handle the microservices. Use Amazon S3 buckets for storage.

B. Use Amazon EC2 instances to host the microservices. Use Amazon Elastic Block Store (Amazon EBS) volumes for storage.

C. Use Amazon Elastic Container Service (Amazon ECS) containers on AWS Fargate to handle the microservices. Use an Amazon Elastic File System (Amazon EFS) file system for storage.

D. Use Amazon Elastic Container Service (Amazon ECS) containers on AWS Fargate to handle the microservices. Use an Amazon EC2 instance that runs a dedicated file store for storage.

Answer: CQUESTION 1295A company is migrating a data processing application to AWS. The application processes several short-lived batch jobs that cannot be disrupted. The process generates data after each batch job finishes running. The company accesses the data for 30 days following data generation. After 30 days, the company stores the data for 2 years. The company wants to optimize costs for the application and data storage. Which solution will meet these requirements?

A. Use Amazon EC2 Spot Instances to run the application. Store the data in Amazon S3 Standard. Move the data to S3 Glacier Instant Retrieval after 30 days. Configure a bucket policy to delete the data after 2 years.

B. Use Amazon EC2 On-Demand Instances to run the application. Store the data in Amazon S3 Glacier Instant Retrieval. Move the data to S3 Glacier Deep Archive after 30 days. Configure an S3 Lifecycle configuration to delete the data after 2 years.

C. Use Amazon EC2 Spot Instances to run the application. Store the data in Amazon S3 Standard. Move the data to S3 Glacier Flexible Retrieval after 30 days. Configure a bucket policy to delete the data after 2 years.

D. Use Amazon EC2 On-Demand Instances to run the application. Store the data in Amazon S3 Standard. Move the data to S3 Glacier Deep Archive after 30 days. Configure an S3 Lifecycle configuration to delete the data after 2 years.

Answer: DQUESTION 1296A genomics research company is designing a scalable architecture for a loosely coupled workload. Tasks in the workload are independent and can be processed in parallel. The architecture needs to minimize management overhead and provide automatic scaling based on demand. Which solution will meet these requirements MOST cost-effectively?

A. Use a cluster of Amazon EC2 instances. Use AWS Systems Manager to manage the workload.

B. Implement a serverless architecture that uses AWS Lambda functions.

C. Use AWS ParallelCluster to deploy a dedicated high-performance cluster.

D. Implement vertical scaling for each workload task.

Answer: BQUESTION 1297A company uses AWS Organizations to manage multiple AWS accounts. Each department in the company has its own AWS account. A security team needs to implement centralized governance and control to enforce security best practices across all accounts. The team wants to have control over which AWS services each account can use. The team needs to restrict access to sensitive resources based on IP addresses or geographic regions. The root user must be protected with multi-factor authentication (MFA) across all accounts. Which solution will meet these requirements?

A. Use AWS Identity and Access Management (IAM) to manage IAM users and IAM roles in each account. Implement MFA for the root user in each account. Enforce service restrictions by using AWS managed prefix lists.

B. Use AWS Control Tower to establish a multi-account environment. Use service control policies (SCPs) to enforce service restrictions in AWS Organizations. Configure MFA for the root user across all accounts.

C. Use AWS Systems Manager to enforce service restrictions across multiple accounts. Use IAM policies to enforce MFA for the root user across all accounts.

D. Use AWS IAM Identity Center to manage user access and to enforce service restrictions by using permissions boundaries in each account.

Answer: BQUESTION 1298An ecommerce company hosts an API that handles sales requests. The company hosts the API frontend on Amazon EC2 instances that run behind an Application Load Balancer (ALB). The company hosts the API backend on EC2 instances that perform the transactions. The backend tiers are loosely coupled by an Amazon Simple Queue Service (Amazon SQS) queue. The company anticipates a significant increase in request volume during a new product launch event. The company wants to ensure that the API can handle increased loads successfully. Which solution will meet these requirements?

A. Double the number of frontend and backend EC2 instances to handle the increased traffic during the product launch event. Create a dead-letter queue to retain unprocessed sales requests when the demand exceeds the system capacity.

B. Place the frontend EC2 instances into an Auto Scaling group. Create an Auto Scaling policy to launch new instances to handle the incoming network traffic.

C. Place the frontend EC2 instances into an Auto Scaling group. Add an Amazon ElastiCache cluster in front of the ALB to reduce the amount of traffic the API needs to handle.

D. Place the frontend and backend EC2 instances into separate Auto Scaling groups. Create a policy for the frontend Auto Scaling group to launch instances based on incoming network traffic. Create a policy for the backend Auto Scaling group to launch instances based on the SQS queue backlog.

Answer: DQUESTION 1299A machine learning (ML) team is building an application that uses data that is in an Amazon S3 bucket. The ML team needs a storage solution for its model training workflow

on AWS. The ML team requires high-performance storage that supports frequent access to training datasets. The storage solution must integrate natively with Amazon S3. Which solution will meet these requirements with the LEAST operational overhead?A. Use Amazon Elastic Block Store (Amazon EBS) volumes to provide high-performance storage. Use AWS DataSync to migrate data from the S3 bucket to EBS volumes.B. Use Amazon EC2 ML instances to provide high-performance storage. Store training data on Amazon EBS volumes. Use the S3 Copy API to copy data from the S3 bucket to EBS volumes.C. Use Amazon FSx for Lustre to provide high-performance storage. Store training datasets in Amazon S3 Standard storage.D. Use Amazon EMR to provide high-performance storage. Store training datasets in Amazon S3 Glacier Instant Retrieval storage. Answer: C

QUESTION 1300A company operates an online photo-sharing service and stores data in AWS Account A in a centralized Amazon S3 bucket. The company wants to grant a second AWS account named Account B access to the centralized S3 bucket. The company owns Account B. Which solution will meet this requirement?A. Enable S3 Transfer Acceleration to provide Account B access to the centralized S3 bucket in Account A.B. Enable cross-Region replication between Account A and Account B to share the S3 bucket data.C. Use Amazon CloudFront to distribute the S3 bucket contents. Grant Account B access to the bucket contents through a signed URL. D. Create a bucket policy that grants Account B permission to access the centralized S3 bucket in Account A. Answer:

D

QUESTION 1301A company wants to migrate an application that uses a microservice architecture to AWS. The services currently run on Docker containers on-premises. The application has an event-driven architecture that uses Apache Kafka. The company configured Kafka to use multiple queues to send and receive messages. Some messages must be processed by multiple services. Which solution will meet these requirements with the LEAST management overhead?A. Migrate the services to Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. Deploy a Kafka cluster on EC2 instances to handle service-to-service communication.B. Migrate the services to Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Create multiple Amazon Simple Queue Service (Amazon SQS) queues to handle service- to-service communication.C. Migrate the services to Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Deploy an Amazon Managed Streaming for Apache Kafka (Amazon MSK) cluster to handle service-to-service communication.D.

Migrate the services to Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. Use Amazon EventBridge to handle service-to-service communication. Answer: C

QUESTION 1302A company wants to migrate an application to AWS. The application runs on Docker containers behind an Application Load Balancer (ALB). The application stores data in a PostgreSQL database. The cloud-based solution must use AWS WAF to inspect all application traffic. The application experiences most traffic on weekdays. There is significantly less traffic on weekends. Which solution will meet these requirements in the MOST cost-effective way? Options:

A. Use a Network Load Balancer (NLB). Create a web access control list (web ACL) in AWS WAF that includes the necessary rules. Attach the web ACL to the NLB. Run the application on Amazon Elastic Container Service (Amazon ECS). Use Amazon RDS for PostgreSQL as the database.B. Create a web access control list (web ACL) in AWS WAF that includes the necessary rules. Attach the web ACL to the ALB. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon RDS for PostgreSQL as the database.C. Create a web access control list (web ACL) in AWS WAF that includes the necessary rules. Attach the web ACL to the ALB. Run the application on Amazon Elastic Container Service (Amazon ECS). Use Amazon Aurora Serverless as the database.D. Use a Network Load Balancer (NLB). Create a web access control list (web ACL) in AWS WAF that has the necessary rules. Attach the web ACL to the NLB. Run the application on Amazon Elastic Container Service (Amazon ECS). Use Amazon Aurora Serverless as the database.

Answer: C

QUESTION 1303A company hosts a public application on AWS. The company uses an Application Load Balancer (ALB) to distribute application traffic to multiple Amazon EC2 instances that are hosted in private subnets. The company wants to authenticate all the requests by using an on-premises Active Directory Federation Service (AD FS). The company uses AWS Direct Connect to connect its on-premises data center to AWS. Which solution will meet this requirement?A. Configure an Amazon Cognito user pool. Integrate the user pool with the ALB for AD FS authentication.B. Configure an AWS Directory Service directory. Integrate the directory with the ALB for AD FS authentication.C. Replace the ALB with a Network Load Balancer (NLB). Use Amazon Connect Agent Workspace to integrate an agent workspace with the NLB.D. Configure an AWS Directory Service AD Connector. Integrate the AD Connector with the ALB for AD FS authentication.

Answer: D

QUESTION 1304A company runs an online order management system on AWS. The company stores order and inventory data for the previous 5 years in an Amazon Aurora MySQL database. The company deletes inventory data after 5 years. The company wants to optimize costs to archive data. Which solution will meet this requirement?A. Create an AWS Glue crawler to export data to Amazon S3. Create an AWS Lambda function to compress the data.B. Use the SELECT INTO OUTFILE S3 query on the Aurora database to export the data to Amazon S3. Configure S3 Lifecycle rules on the S3 bucket.C. Create an AWS Glue DataBrew Job to migrate data from Aurora to Amazon S3. Configure S3 Lifecycle rules on the S3 bucket.D. Use the AWS Schema Conversion Tool (AWS SCT) to replicate data from Aurora to Amazon S3. Use the S3

Standard-Infrequent Access (S3 Standard-IA) storage class. Answer: BQUESTION 1305A company runs a mobile game app that stores session data (up to 256 KB) for up to 48 hours. The data updates frequently and must be deleted automatically after expiration. Restorability is also required. Which solution will meet this requirement? A. Use an Amazon DynamoDB table to store the session data. Enable point-in-time recovery (PITR) and TTL. B. Use Amazon MemoryDB and enable PITR and TTL. C. Store session data in S3 Standard. Enable Versioning and a Lifecycle rule to expire objects after 48 hours. D. Store data in S3 Intelligent-Tiering with Versioning and a Lifecycle rule to expire after 48 hours. Answer: AQUESTION 1306A finance company collects streaming data for a real-time search and visualization system. They want to migrate to AWS using a native solution for ingest, search, and visualization. Which solution will meet this requirement? A. Use EC2 to ingest/process data to S3 Athena + Managed Grafana. B. Use EMR to ingest/process to Redshift Redshift Spectrum + QuickSight. C. Use EKS to ingest/process to DynamoDB CloudWatch Dashboards. D. Use Kinesis Data Streams Amazon OpenSearch Service Amazon QuickSight. Answer: DQUESTION 1307A company uses Apache Hadoop and Spark on-prem. The infrastructure is complex and not scalable. They want to reduce operational complexity but keep data processing on-premises. Which solution will meet this requirement? A. Use Site-to-Site VPN to access on-prem HDFS. Use Amazon EMR to process the data. B. Use AWS DataSync to connect to on-prem HDFS. Use Amazon EMR to process the data. C. Migrate to Amazon EMR on AWS Outposts. D. Use AWS Snowball to migrate data to S3. Use EMR to process. Answer: CQUESTION 1308A company recently migrated a large amount of research data to an Amazon S3 bucket. The company needs an automated solution to identify sensitive data in the bucket. A security team also needs to monitor access patterns for the data 24 hours a day, 7 days a week to identify suspicious activities or evidence of tampering with security controls. Which solution will meet this requirement? A. Set up AWS CloudTrail reporting, and grant the security team read-only access to the CloudTrail reports. Set up an Amazon S3 Inventory report to identify sensitive data. Review the findings with the security team. B. Enable Amazon Macie and Amazon GuardDuty on the account. Grant the security team access to Macie and GuardDuty. Review the findings with the security team. C. Set up an Amazon S3 Inventory report. Use Amazon Athena and Amazon QuickSight to identify sensitive data. Create a dashboard for the security team to review findings. D. Use AWS Identity and Access Management (IAM) Access Advisor to monitor for suspicious activity and tampering. Create a dashboard for the security team. Set up an Amazon S3 Inventory report to identify sensitive data. Review the findings with the security team. Answer: BQUESTION 1309A healthcare company uses an Amazon EMR cluster to process patient data. The data must be encrypted in transit and at rest. Local volumes in the cluster also need to be encrypted. Which solution will meet these requirements? A. Create Amazon EBS volumes. Enable encryption. Attach the volumes to the existing EMR cluster. B. Create an EMR security configuration that encrypts the data and the volumes as required. C. Create an EC2 instance profile for the EMR instances. Configure the instance profile to enforce encryption. D. Create a runtime role that has a trust policy for the EMR cluster. Answer: BQUESTION 1310A company is building an ecommerce application that uses a relational database to store customer data and order history. The company also needs a solution to store 100 GB of product images. The company expects the traffic flow for the application to be predictable. Which solution will meet these requirements MOST cost-effectively? A. Use Amazon RDS for MySQL for the database. Store the product images in an Amazon S3 bucket. B. Use Amazon DynamoDB for the database. Store the product images in an Amazon S3 bucket. C. Use Amazon RDS for MySQL for the database. Store the product images in an Amazon Aurora MySQL database. D. Create three Amazon EC2 instances. Install MongoDB software on the instances to use as the database. Store the product images in an Amazon RDS for MySQL database with a Multi-AZ deployment. Answer: AQUESTION 1311A company runs a monolithic application in its on-premises data center. The company used Java/Tomcat to build the application. The application uses Microsoft SQL Server as a database. The company wants to migrate the application to AWS. Which solution will meet this requirement with the LEAST operational overhead? A. Use AWS App2Container to containerize the application. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS). Deploy the database to Amazon RDS for SQL Server. Configure a Multi-AZ deployment. B. Containerize the application and deploy the application on a self-managed Kubernetes cluster on an Amazon EC2 instance. Deploy the database on a separate EC2 instance. Set up Microsoft SQL Server Always On availability groups. C. Deploy the frontend of the web application as a website on Amazon S3. Use Amazon DynamoDB for the database tier. D. Use AWS App2Container to containerize the application. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon DynamoDB for the database tier. Answer: AQUESTION 1312A company uses a single Amazon S3 bucket to store data that multiple business applications must access. The company hosts the applications on Amazon EC2 Windows instances that are in a VPC. The company configured a bucket policy for the S3 bucket to grant the applications access to the bucket. The company continually adds more business applications to the environment. As the number of business applications increases, the policy document becomes more difficult to manage. The S3 bucket policy document will soon reach its policy size quota. The company needs a solution to scale its architecture to handle more business applications. Which

solution will meet these requirements in the MOST operationally efficient way?A. Migrate the data from the S3 bucket to an Amazon Elastic File System (Amazon EFS) volume. Ensure that all application owners configure their applications to use the EFS volume.B. Deploy an AWS Storage Gateway appliance for each application. Reconfigure the applications to use a dedicated Storage Gateway appliance to access the S3 objects instead of accessing the objects directly.C. Create a new S3 bucket for each application. Configure S3 replication to keep the new buckets synchronized with the original S3 bucket. Instruct application owners to use their respective S3 buckets.D. Create an S3 access point for each application. Instruct application owners to use their respective S3 access points.

Answer: D

QUESTION 1313A company needs to store confidential files on AWS. The company accesses the files every week. The company must encrypt the files by using envelope encryption, and the encryption keys must be rotated automatically. The company must have an audit trail to monitor encryption key usage. Which combination of solutions will meet these requirements? (Choose two.)

A. Store the confidential files in Amazon S3.B. Store the confidential files in Amazon S3 Glacier Deep Archive.C. Use server-side encryption with customer-provided keys (SSE-C).D. Use server-side encryption with Amazon S3 managed keys (SSE-S3).E. Use server-side encryption with AWS KMS managed keys (SSE-KMS).

Answer: AE

QUESTION 1314An ecommerce company is launching a new marketing campaign. The company anticipates the campaign to generate ten times the normal number of daily orders through the company's ecommerce application. The campaign will last 3 days. The ecommerce application architecture is based on Amazon EC2 instances in an Auto Scaling group and an Amazon RDS for MySQL database. The application writes order transactions to an Amazon Elastic File System (Amazon EFS) file system before the application writes orders to the database. During normal operations, the application write operations peak at 5,000 IOPS. A solutions architect needs to ensure that the application can handle the anticipated workload during the marketing campaign. Which solution will meet this requirement?

A. For the duration of the campaign, increase the provisioned IOPS for the RDS for MySQL database. Set the Amazon EFS throughput mode to Bursting throughput.B. For the duration of the campaign, increase the provisioned IOPS for the RDS for MySQL database. Set the Amazon EFS throughput mode to Elastic throughput.C. Convert the database to a Multi-AZ deployment. Set the Amazon EFS throughput mode to Elastic throughput for the duration of the campaign.D. Use AWS Database Migration Service (AWS DMS) to convert the database to RDS for PostgreSQL. Set the Amazon EFS throughput mode to Bursting throughput.

Answer: B

QUESTION 1315A company runs an online order management system on AWS. The company stores order and inventory data for the previous 5 years in an Amazon Aurora MySQL database. The company deletes inventory data after 5 years. The company wants to optimize costs to archive data. Which solution will meet this requirement?

A. Create an AWS Glue crawler to export data to Amazon S3. Create an AWS Lambda function to compress the data.B. Use the SELECT INTO OUTFILE S3 query on the Aurora database to export the data to Amazon S3. Configure S3 Lifecycle rules on the S3 bucket.C. Create an AWS Glue DataBrew job to migrate data from Aurora to Amazon S3. Configure S3 Lifecycle rules on the S3 bucket.D. Use the AWS Schema Conversion Tool (AWS SCT) to replicate data from Aurora to Amazon S3. Use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.

Answer: B

QUESTION 1316A company is designing an advertisement distribution application to run on AWS. The company wants to deploy the application as a container to Amazon Elastic Container Service (Amazon ECS). Advertisements must be displayed to users around the world with low latency. The company needs to optimize data transfer costs. Which solution will meet these requirements?

A. Deploy the application in a single AWS Region. Use an Application Load Balancer (ALB) to distribute traffic. Create an Amazon CloudFront distribution, and set the ALB as the origin.B. Deploy the application in multiple AWS Regions. Create an Application Load Balancer (ALB) in each Region. Use Amazon Route 53 with a latency-based weighted routing policy to distribute traffic to the ALBs.C. Deploy the application in multiple AWS Regions. Create an Application Load Balancer (ALB) in each Region. Create a transit gateway in each Region. Route traffic between the ALBs and Amazon ECS through the transit gateways.D. Deploy the application in a single AWS Region. Use an Application Load Balancer (ALB) to distribute traffic. Create an accelerator in AWS Global Accelerator. Associate the accelerator with the ALB.

Answer: A

QUESTION 1317A company runs a three-tier web application in a VPC on AWS. The company deployed an Application Load Balancer (ALB) in a public subnet. The web tier and application tier Amazon EC2 instances are deployed in a private subnet. The company uses a self-managed MySQL database that runs on EC2 instances in an isolated private subnet for the database tier. The company wants a mechanism that will give a DevOps team the ability to use SSH to access all the servers. The company also wants to have a centrally managed log of all connections made to the servers. Which combination of solutions will meet these requirements with the MOST operational efficiency? (Choose two.)

A. Create a bastion host in the public subnet. Configure security groups in the public, private, and isolated subnets to allow SSH access.B. Create an interface VPC endpoint for AWS Systems Manager Session Manager. Attach the endpoint to the VPC.C. Create an IAM policy that grants access to AWS Systems Manager Session Manager. Attach the IAM policy to the EC2 instances.D. Create a gateway VPC endpoint for AWS Systems Manager Session Manager. Attach the endpoint to the VPC.E. Attach an AmazonSSMManagedInstanceCore AWS managed IAM policy to all the

EC2 instance roles. Answer: BEQUESTION 1318A company needs to grant a team of developers access to the company's AWS resources. The company must maintain a high level of security for the resources. The company requires an access control solution that will prevent unauthorized access to the sensitive data. Which solution will meet these requirements?A. Share the IAM user credentials for each development team member with the rest of the team to simplify access management and to streamline development workflows.B. Define IAM roles that have fine-grained permissions based on the principle of least privilege. Assign an IAM role to each developer.C. Create IAM access keys to grant programmatic access to AWS resources. Allow only developers to interact with AWS resources through API calls by using the access keys.D. Create an AWS Cognito user pool. Grant developers access to AWS resources by using the user pool. Answer: BQUESTION 1319A company is developing a serverless, bidirectional chat application that can broadcast messages to connected clients. The application is based on AWS Lambda functions. The Lambda functions receive incoming messages in JSON format. The company needs to provide a frontend component for the application. Which solution will meet this requirement?A. Use an Amazon API Gateway HTTP API to direct incoming JSON messages to backend destinations.B. Use an Amazon API Gateway REST API that is configured with a Lambda proxy integration.C. Use an Amazon API Gateway WebSocket API to direct incoming JSON messages to backend destinations.D. Use an Amazon CloudFront distribution that is configured with a Lambda function URL as a custom origin. Answer: CQUESTION 1320A company is setting up a development environment on AWS for a team of developers. The team needs to access multiple Amazon S3 buckets to store project data. The team also needs to use Amazon EC2 to run development instances. The company needs to ensure that the developers have access only to specific Amazon S3 buckets and EC2 instances. Access permissions must be assigned according to each developer's role on the team. The company wants to minimize the use of permanent credentials and to ensure access is securely managed according to the principle of least privilege. Which solution will meet these requirements?A. Create IAM roles that have administrative-level permissions for Amazon S3 and Amazon EC2. Require developers to sign in by using Amazon Cognito to access Amazon S3 and Amazon EC2.B. Create IAM roles that have fine-grained permissions for Amazon S3 and Amazon EC2. Configure AWS IAM Identity Center to manage credentials for the developers.C. Create IAM users that have programmatic access to Amazon S3 and Amazon EC2. Generate individual access keys for each developer to access Amazon S3 and Amazon EC2.D. Create a VPC endpoint for Amazon S3. Require developers to access Amazon EC2 instances and Amazon S3 buckets through a bastion host. Answer: BQUESTION 1321A company uses an Amazon EC2 instance to handle requests for a public web application. The application routes traffic to multiple application pages by using URL paths. The company begins to experience large surges of traffic at unpredictable times. The traffic surges cause the web application to experience issues and to occasionally become unavailable. The company needs to make the web application more scalable to handle sudden increases in traffic. Which solution will meet this requirement?A. Create an Amazon Machine Image (AMI) of the web application instance. Use the AMI to create an Auto Scaling group of EC2 instances that has a minimum capacity of two. Create an Application Load Balancer. Set the Auto Scaling group as the target group.B. Create a Docker image of the application. Use Amazon Elastic Container Service (Amazon ECS) to create an Auto Scaling ECS cluster. Enable managed scaling. Create a Network Load Balancer. Set the ECS cluster as the target group.C. Create an Amazon Machine Image (AMI) of the web application instance. Use the AMI to create two more web application instances in separate Availability Zones. Update the website DNS record to refer to all three instances.D. Create an Application Load Balancer (ALB). Set the web application instance as the target. Create an Amazon CloudWatch alarm based on ALB traffic metrics. Configure the alert to activate when traffic spikes. Answer: AQUESTION 1322A company is using an Amazon Redshift cluster to run analytics queries for multiple sales teams. In addition to the typical workload, on the last Monday morning of each month, thousands of users run reports. Users have reported slow response times during the monthly surge. The company must improve query performance without impacting the availability of the Redshift cluster. Which solution will meet these requirements?A. Resize the Redshift cluster by using the classic resize capability of Amazon Redshift before every monthly surge. Reduce the cluster to its original size after each surge.B. Resize the Redshift cluster by using the elastic resize capability of Amazon Redshift before every monthly surge. Reduce the cluster to its original size after each surge.C. Enable the concurrency scaling feature for the Redshift cluster for specific workload management (WLM) queues.D. Enable Amazon Redshift Spectrum for the Redshift cluster before every monthly surge. Answer: CQUESTION 1323A company hosts a web application in a VPC on AWS. A public Application Load Balancer (ALB) forwards connections from the internet to an Auto Scaling group of Amazon EC2 instances. The Auto Scaling group runs in private subnets across four Availability Zones. The company stores data in an Amazon S3 bucket in the same Region. The EC2 instances use NAT gateways in each Availability Zone for outbound internet connectivity. The company wants to optimize costs for its AWS architecture. Which solution will meet this requirement?A. Reconfigure the Auto Scaling group and the ALB to use two Availability Zones instead of four. Do not change the desired count or scaling metrics for the Auto Scaling group to maintain application availability.B. Create a new, smaller VPC that still has sufficient IP address

availability to run the application. Redeploy the application stack in the new VPC. Delete the existing VPC and its resources.C. Deploy an S3 gateway endpoint to the VPC. Configure the EC2 instances to access the S3 bucket through the S3 gateway endpoint.

D. Deploy an S3 interface endpoint to the VPC. Configure the EC2 instances to access the S3 bucket through the S3 interface endpoint.

Answer: C

QUESTION 1324A company needs to run its external website on Amazon EC2 instances and on-premises virtualized servers. The AWS environment has a 1 GB AWS Direct Connect connection to the data center. The application has IP addresses that will not change. The on-premises and AWS servers are able to restart themselves while maintaining the same IP address if a failure occurs. Some website users have to add their vendors to an allow list, so the solution must have a fixed IP address. The company needs a solution with the lowest operational overhead to handle this split traffic. What should a solutions architect do to meet these requirements?

A. Deploy an Amazon Route 53 Resolver with rules pointing to the on-premises and AWS IP addresses.

B. Deploy a Network Load Balancer on AWS. Create target groups for the on-premises and AWS IP addresses.

C. Deploy an Application Load Balancer on AWS. Register the on-premises and AWS IP addresses with the target group.

D. Deploy Amazon API Gateway to direct traffic to the on-premises and AWS IP addresses based on the header of the request.

Answer: B

QUESTION 1325A company runs multiple applications in multiple AWS accounts within the same organization in AWS Organizations. A content management system (CMS) runs on Amazon EC2 instances in a VPC. The CMS needs to access shared files from an Amazon Elastic File System (Amazon EFS) file system that is deployed in a separate AWS account. The EFS account is in a separate VPC. Which solution will meet this requirement?

A. Mount the EFS file system on the EC2 instances by using the EFS Elastic IP address.

B. Enable VPC sharing between the two accounts. Use the EFS mount helper to mount the file system on the EC2 instances. Redeploy the EFS file system in a shared subnet.

C. Configure AWS Systems Manager Run Command to mount the EFS file system on the EC2 instances.

D. Install the amazon-efs-utils package on the EC2 instances. Add the mount target in the efs-config file. Mount the EFS file system by using the EFS access point.

Answer: D

QUESTION 1326A company runs a web application that uses Amazon RDS for MySQL to store relational data. Data in the database does not change frequently. A solutions architect notices that during peak usage times, the database has performance issues when it serves the data. The company wants to improve the performance of the database. Which combination of steps will meet these requirements? (Choose two.)

A. Integrate AWS WAF with the application.

B. Create a read replica for the database. Redirect read traffic to the read replica.

C. Create an Amazon ElastiCache (Memcached) cluster. Configure the application and the database to integrate with the cluster.

D. Use the Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class to store the data that changes infrequently.

E. Migrate the database to Amazon DynamoDB. Configure the application to use the DynamoDB database.

Answer: B

QUESTION 1327A company receives data transfers from a small number of external clients that use SFTP software on an Amazon EC2 instance. The clients use an SFTP client to upload data. The clients use SSH keys for authentication. Every hour, an automated script transfers new uploads to an Amazon S3 bucket for processing. The company wants to move the transfer process to an AWS managed service and to reduce the time required to start data processing. The company wants to retain the existing user management and SSH key generation process. The solution must not require clients to make significant changes to their existing processes. Which solution will meet these requirements?

A. Reconfigure the script that runs on the EC2 instance to run every 15 minutes. Create an S3 Event Notifications rule for all new object creation events. Set an Amazon Simple Notification Service (Amazon SNS) topic as the destination.

B. Create an AWS Transfer Family SFTP server that uses the existing S3 bucket as a target. Use service-managed users to enable authentication.

C. Require clients to add the AWS DataSync agent into their local environments. Create an IAM user for each client that has permission to upload data to the target S3 bucket.

D. Create an AWS Transfer Family SFTP connector that has permission to access the target S3 bucket for each client. Store credentials in AWS Systems Manager. Create an IAM role to allow the SFTP connector to securely use the credentials.

Answer: B

QUESTION 1328A company needs to run a critical data processing workload that uses a Python script every night. The workload takes 1 hour to finish. Which solution will meet these requirements MOST cost-effectively?

A. Deploy an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type. Use the Fargate Spot capacity provider. Schedule the job to run once every night.

B. Deploy an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type. Schedule the job to run once every night.

C. Create an AWS Lambda function that uses the existing Python code. Configure Amazon EventBridge to invoke the function once every night.

D. Create an Amazon EC2 On-Demand Instance that runs Amazon Linux. Migrate the Python script to the instance. Use a cron job to schedule the script. Create an AWS Lambda function to start and stop the instance once every night.

Answer: C

QUESTION 1329A solutions architect is creating a data reporting application that will send traffic through third-party network firewalls in an AWS security account. The firewalls and application servers must be load balanced. The application uses TCP connections to generate reports. The reports can run for several hours and can be idle for up to 1 hour. The reports must not time out during an idle period. Which solution will meet these requirements?

A. Use a Gateway Load Balancer (GWLB) for the firewalls.

Use an Application Load Balancer (ALB) for the application servers. Set the ALB idle timeout period to 1 hour.B. Use a single firewall in the security account. Use an Application Load Balancer (ALB) for the application servers. Set the ALB idle timeout and firewall idle timeout periods to 1 hour.C. Use a Gateway Load Balancer (GWLB) for the firewalls. Use an Application Load Balancer (ALB) for the application servers. Set the idle timeout periods for the ALB, the GWLB, and the firewalls to 1 hour.D. Use a Gateway Load Balancer (GWLB) for the firewalls. Use an Application Load Balancer (ALB) for the application servers. Configure the ALB idle timeout period to 1 hour. Increase the application server capacity to finish the report generation faster.

Answer: C

QUESTION 1330A company collects 10 GB of telemetry data every day from multiple devices. The company stores the data in an Amazon S3 bucket that is in a source data account. The company has hired several consulting agencies to analyze the company's data. Each agency has a unique AWS account. Each agency requires read access to the company's data. The company needs a secure solution to share the data from the source data account to the consulting agencies. Which solution will meet these requirements with the LEAST operational effort?

A. Set up an Amazon CloudFront distribution. Use the S3 bucket as the origin.

B. Make the S3 bucket public for a limited time. Inform only the agencies that the bucket is publicly accessible.

C. Configure cross-account access for the S3 bucket to the accounts that the agencies own.

D. Set up an IAM user for each agency in the source data account. Grant each agency IAM user access to the company's S3 bucket.

Answer: C

QUESTION 1331A company is migrating some workloads to AWS. However, many workloads will remain on premises. The on-premises workloads require secure and reliable connectivity to AWS with consistent, low- latency performance. The company has deployed the AWS workloads across multiple AWS accounts and multiple VPCs. The company plans to scale to hundreds of VPCs within the next year. The company must establish connectivity between each of the VPCs and from the on-premises environment to each VPC. Which solution will meet these requirements?

A. Use an AWS Direct Connect connection to connect the on-premises environment to AWS.

B. Configure VPC peering to establish connectivity between VPCs.

C. Use an AWS Direct Connect connection with a Direct Connect gateway to connect the on-premises environment to AWS. Create a transit gateway to establish connectivity between VPCs.

D. Use an AWS Site-to-Site VPN connection to connect the on-premises environment to AWS.

Answer: C

Resources From: 1.2025 Latest Braindump2go SAA-C03 Exam Dumps (PDF & VCE) Free Share:

<https://www.braindump2go.com/saa-c03.html> 2.025 Latest Braindump2go SAA-C03 PDF and SAA-C03 VCE Dumps Free

Share: [https://drive.google.com/drive/folders/1PKc\\_AsNW5xtYjJaY4\\_oJFcTRLkRk9lPW?usp=sharing](https://drive.google.com/drive/folders/1PKc_AsNW5xtYjJaY4_oJFcTRLkRk9lPW?usp=sharing) 3.2025 Free

Braindump2go SAA-C03 Exam Questions Download:

[https://www.braindump2go.com/free-online-pdf/SAA-C03-VCE-Dumps\(976-1010\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C03-VCE-Dumps(976-1010).pdf) Free Resources from

Braindump2go, We Devoted to Helping You 100% Pass All Exams!