Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

https://www.2passeasy.com/dumps/SAP-C02/

NEW QUESTION 1
- (Exam Topic 1)
A solutions architect is designing a publicly accessible web application that is on an Amazon CloudFront distribution with an Amazon S3 website endpoint as the origin. When the solution is deployed, the website returns an Error 403: Access Denied message.
Which steps should the solutions architect take to correct the issue? (Select TWO.)

A. Remove the S3 block public access option from the S3 bucket.
B. Remove the requester pays option trom the S3 bucket.
C. Remove the origin access identity (OAI) from the CloudFront distribution.
D. Change the storage class from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA).
E. Disable S3 object versioning.

Answer: AB

Explanation:
See using S3 to host a static website with Cloudfront: https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-serve-static-website/
- Using a REST API endpoint as the origin, with access restricted by an origin access identity (OAI)
- Using a website endpoint as the origin, with anonymous (public) access allowed
- Using a website endpoint as the origin, with access restricted by a Referer header

NEW QUESTION 2
- (Exam Topic 1)
A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.
The company has the following DNS resolution requirements:
• On-premises systems should be able to resolve and connect to cloud.example.com.
• All VPCs should be able to resolve cloud.example.com.
There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway. Which architecture should the company use to meet these requirements with the HIGHEST performance?

A. Associate the private hosted zone to all the VPC
B. Create a Route 53 inbound resolver in the shared services VP
C. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.
D. Associate the private hosted zone to all the VPC
E. Deploy an Amazon EC2 conditional forwarder in the shared services VP
F. Attach all VPCs to the transit gateway and create forwarding rules in theon-premises DNS server for cloud.example.com that point to the conditional forwarder.
G. Associate the private hosted zone to the shared services VP
H. Create a Route 53 outbound resolver in the shared services VP
I. Attach all VPCs to the transit gateway and create forwarding rules in theon-premises DNS server for cloud.example.com that point to the outbound resolver.
J. Associate the private hosted zone to the shared services VP
K. Create a Route 53 inbound resolver in the shared services VP
L. Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

Answer: D

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-w

NEW QUESTION 3
- (Exam Topic 1)
An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory requirement for out-of-region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles.
Which of the following options can the solutions architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

A. Back up the application and database data frequently and copy them to Amazon S3. Replicate the backups using S3 cross-region replication, and use AWS Cloud Formation to instantiate infrastructure for disaster recovery and restore data from Amazon S3.
B. Employ a pilot light environment in which the primary database is configured with mirroring to build a standby database on m4.large in Ihe alternate regio
C. Use AWS Cloud Formation to instantiate the web servers, application servers, and load balancers in case of a disaster to bring the application up in the alternate regio
D. Vertically resize the database to meet the full production demands, and use Amazon Route 53 to switch traffic to the alternate region.
E. Use a scaled-down version of the fully functional production environment in the alternate region that includes one instance of the web server, one instance of the application server, and a replicated instance of the database server in standby mod
F. Place the web and the application tiers in an Auto Scaling group behind a load balancer, which can automatically scale when the load arrives to the applicatio
G. Use Amazon Route 53 to switch traffic to the alternate region,
H. Employ a multi-region solution with fully functional we
I. application, and database tiers in both regions with equivalent capacit
J. Activate the primary database in one region only and the standby database in the other regio
K. Use Amazon Route 53 to automatically switch traffic from one region to another using health check routing policies.

Answer: C

Explanation:
As RTO is in minutes
(https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/plan-for-disaster-recovery-dr.html) Warm standby (RPO in seconds, RTO in minutes): Maintain a scaled-down version of a fully functional environment always running in the DR Region. Business-critical systems are fully duplicated and are always on, but with

a scaled down fleet. When the time comes for recovery, the system is scaled up quickly to handle the production load.

NEW QUESTION 4
- (Exam Topic 1)
A development team has created a new flight tracker application that provides near-real-time data to users. The application has a front end that consists of an Application Load Balancer (ALB) in front of two large Amazon EC2 instances in a single Availability Zone. Data is stored in a single Amazon RDS MySQL DB instance. An Amazon Route 53 DNS record points to the ALB.
Management wants the development team to improve the solution to achieve maximum reliability with the least amount of operational overhead.
Which set of actions should the team take?

A. Create RDS MySQL read replica
B. Deploy the application to multiple AWS Region
C. Use a Route 53 latency-based routing policy to route to the application.
D. Configure the DB instance as Multi-A
E. Deploy the application to two additional EC2 instances in different Availability Zones behind an ALB.
F. Replace the DB instance with Amazon DynamoDB global table
G. Deploy the application in multiple AWS Region
H. Use a Route 53 latency-based routing policy to route to the application.
I. Replace the DB instance with Amazon Aurora with Aurora Replica
J. Deploy the application to mulliple smaller EC2 instances across multiple Availability Zones in an Auto Scaling group behind an ALB.

Answer: D

Explanation:
Multi AZ ASG + ALB + Aurora = Less over head and automatic scaling

NEW QUESTION 5
- (Exam Topic 1)
A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.
The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.
Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?

A. Deploy a Lambda@Edge function to sort parameters by name and force them to be lowercas
B. Select the CloudFront viewer request trigger to invoke the function.
C. Update the CloudFront distribution to disable caching based on query string parameters.
D. Deploy a reverse proxy after the load balancer to post-process the emitted URLs in the application to force the URL strings to be lowercase.
E. Update the CloudFront distribution to specify casing-insensitive query string processing.

Answer: A

Explanation:
https://docs.amazonaws.cn/en_us/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-ex Before CloudFront serves content from the cache it will trigger any Lambda function associated with the Viewer Request, in which we can normalize parameters.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-examp

NEW QUESTION 6
- (Exam Topic 1)
A company has an application that sells tickets online and experiences bursts of demand every 7 days. The application has a stateless presentation layer running on Amazon EC2. an Oracle database to store unstructured data catalog information, and a backend API layer. The front-end layer uses an Elastic Load Balancer to distribute the load across nine On-Demand Instances over three Availability Zones (AZs). The Oracle database is running on a single EC2 instance. The company is experiencing performance issues when running more than two concurrent campaigns. A solutions architect must design a solution that meets the following requirements:
• Address scalability issues.
• Increase the level of concurrency.
• Eliminate licensing costs.
• Improve reliability.
Which set of steps should the solutions architect take?

A. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
B. Convert the Oracle database into a single Amazon RDS reserved DB instance.
C. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
D. Create two additional copies of the database instance, then distribute the databases in separate AZs.
E. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
F. Convert the tables in the Oracle database into Amazon DynamoDB tables.
G. Convert the On-Demand Instances into Spot Instances to reduce costs for the front en
H. Convert the tables in the Oracle database into Amazon DynamoDB tables.

Answer: C

Explanation:
Combination of On-Demand and Spot Instances + DynamoDB.

NEW QUESTION 7
- (Exam Topic 1)
A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.
Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed Administrators also must have the ability to

automatically update and remediate noncompliant AWS WAF rules in all accounts
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organizatio
B. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage Update the parameter as needed to add or remove accounts or OUs Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account
C. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rule
D. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resources Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
E. Create AWS WAF rules in the management account of the organization Use AWS Lambda environment variables to store account numbers and OUs to manage Update environment variables as needed to add or remove accounts or OUs Create cross-account IAM roles in member accounts Assume the rotes by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts.
F. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage Update AWS KMS as needed to add or remove accounts or OUs Create IAM users in member accounts Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: B

NEW QUESTION 8
- (Exam Topic 1)
A travel company built a web application that uses Amazon Simple Email Service (Amazon SES) to send email notifications to users. The company needs to enable logging to help troubleshoot email delivery issues. The company also needs the ability to do searches that are based on recipient, subject, and time sent. Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

A. Create an Amazon SES configuration set with Amazon Kinesis Data Firehose as the destinatio
B. Choose to send logs to an Amazon S3 bucket.
C. Enable AWS CloudTrail loggin
D. Specify an Amazon S3 bucket as the destination for the logs.
E. Use Amazon Athena to query the fogs in the Amazon S3 bucket for recipient, subject, and time sent.
F. Create an Amazon CloudWatch log grou
G. Configure Amazon SES to send logs to the log group
H. Use Amazon Athena to query the logs in Amazon CloudWatch for recipient, subject, and time sent.

Answer: AC

Explanation:
https://docs.aws.amazon.com/ses/latest/dg/event-publishing-retrieving-firehose.html
To enable you to track your email sending at a granular level, you can set up Amazon SES to publish email sending events to Amazon CloudWatch, Amazon Kinesis Data Firehose, or Amazon Simple Notification Service based on characteristics that you define.
https://docs.aws.amazon.com/ses/latest/dg/monitor-using-event-publishing.html
https://aws.amazon.com/getting-started/hands-on/build-serverless-real-time-data-processing-app-lambda-kinesis

NEW QUESTION 9
- (Exam Topic 1)
A finance company is running its business-critical application on current-generation Linux EC2 instances The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.
Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.
B. Performing a one-time migration of the database cluster to Amazon RD
C. and creating several additional read replicas to handle the load during end of month
D. Using Amazon CioudWatch with AWS Lambda to change the typ
E. size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric
F. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Answer: B

Explanation:
In this scenario, the Amazon EC2 instances are in an Auto Scaling group already which means that the database read operations is the possible bottleneck especially during the month-end wherein the reports are generated. This can be solved by creating RDS read replicas.

NEW QUESTION 10
- (Exam Topic 1)
A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.
All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX), and must accommodate high levels of throughput.
Which storage solution will meet these requirements?

A. Provision an AWS Storage Gateway file gateway NFS file share that is attached to an Amazon S3 bucke
B. Mount the NFS file share on each EC2 instance In the cluster.
C. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses General Purpose performance mod
D. Mount the EFS file system on each EC2 instance in the cluster.
E. Provision a new Amazon Elastic Block Store (Amazon EBS) volume that uses the lo2 volume type.Attach the EBS volume to all of the EC2 instances in the cluster.
F. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mod

G. Mount the EFS file system on each EC2 instance in the cluster.

Answer: D

NEW QUESTION 10
- (Exam Topic 1)
A company wants to host a new global website that consists of static content. A solutions architect is working on a solution that uses Amazon CloudFront with an origin access identity <OAI) to access website content that is stored in a private Amazon S3 bucket.
During testing, the solutions architect receives 404 errors from the S3 bucket. Error messages appear only for attempts to access paths that end with a forward slash. such as example.com/path/. These requests should return the existing S3 object path/index.html. Any potential solution must not prevent CloudFront from caching the content.
What should the solutions architect do to resolve this problem?

A. Change the CloudFront origin to an Amazon API Gateway proxy endpoin
B. Rewrite the S3 request URL by using an AWS Lambda function.
C. Change the CloudFront origin to an Amazon API Gateway endpoin
D. Rewrite the S3 request URL in an AWS service integration.
E. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by a viewer request event to rewrite the S3 request URL.
F. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by an origin request event to rewrite the S3 request URL.

Answer: C

NEW QUESTION 12
- (Exam Topic 1)
A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.
A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.
Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage clas
B. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loadin
C. Use the new file system as the shared storage for the duration of the jo
D. Delete the file system when the job is complete.
E. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enable
F. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch templat
G. Use the EBS volume as the shared storage for the duration of the jo
H. Detach the EBS volume when the job is complete.
I. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage clas
J. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loadin
K. Use the new file system as the shared storage for the duration of the jo
L. Delete the file system when the job is complete.
M. Migrate the data from the existing shared file system to an Amazon S3 bucke
N. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the jo
O. Delete the file gateway when the job is complete.

Answer: B

NEW QUESTION 13
- (Exam Topic 1)
An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a solutions architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible (or receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.
Which of the following is the MOST reliable approach to meet the requirements?

A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.
B. Receive the orders in an Amazon SOS queue and trigger an AWS Lambda function lo process them.
C. Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container lo process them.
D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

Answer: B

Explanation:
Q: How does Amazon Kinesis Data Streams differ from Amazon SQS?
Amazon Kinesis Data Streams enables real-time processing of streaming big data. It provides ordering of records, as well as the ability to read and/or replay records in the same order to multiple Amazon Kinesis Applications. The Amazon Kinesis Client Library (KCL) delivers all records for a given partition key to the same record processor, making it easier to build multiple applications reading from the same Amazon Kinesis data stream (for example, to perform counting, aggregation, and filtering).
https://aws.amazon.com/kinesis/data-streams/faqs/
https://aws.amazon.com/blogs/big-data/unite-real-time-and-batch-analytics-using-the-big-data-lambda-architect

NEW QUESTION 16
- (Exam Topic 1)
A company standardized its method of deploying applications to AWS using AWS CodePipeline and AWS Cloud Formation. The applications are in Typescript and Python. The company has recently acquired another business that deploys applications to AWS using Python scripts.

Developers from the newly acquired company are hesitant to move their applications under CloudFormation because it would require than they learn a new domain-specific language and eliminate their access to language features, such as looping.
How can the acquired applications quickly be brought up to deployment standards while addressing the developers' concerns?

A. Create CloudFormation templates and re-use parts of the Python scripts as instance user dat
B. Use the AWS Cloud Development Kit (AWS CDK) to deploy the application using these template
C. Incorporate the AWS CDK into CodePipeline and deploy the application to AWS using these templates.
D. Use a third-party resource provisioning engine inside AWS CodeBuild to standardize the deployment processes of the existing and acquired compan
E. Orchestrate the CodeBuild job using CodePipeline.
F. Standardize on AWS OpsWork
G. Integrate OpsWorks with CodePipelin
H. Have the developers create Chef recipes to deploy their applications on AWS.
I. Define the AWS resources using Typescript or Pytho
J. Use the AWS Cloud Development Kit (AWS CDK) to create CloudFormation templates from the developers' code, and use the AWS CDK to create CloudFormation stack
K. Incorporate the AWS CDK as a CodeBuild job in CodePipeline.

Answer: D

NEW QUESTION 20
- (Exam Topic 1)
A company is using AWS Organizations lo manage multiple accounts. Due to regulatory requirements, the company wants to restrict specific member accounts to certain AWS Regions, where they are permitted to deploy resources. The resources in the accounts must be tagged, enforced based on a group standard, and centrally managed with minimal configuration.
What should a solutions architect do to meet these requirements?

A. Create an AWS Config rule in the specific member accounts to limit Regions and apply a tag policy.
B. From the AWS Billing and Cost Management console, in the master account, disable Regions for the specific member accounts and apply a tag policy on the root.
C. Associate the specific member accounts with the roo
D. Apply a tag policy and an SCP using conditions to limit Regions.
E. Associate the specific member accounts with a new O
F. Apply a tag policy and an SCP using conditions to limit Regions.

Answer: D

NEW QUESTION 24
- (Exam Topic 1)
A web application is hosted in a dedicated VPC that is connected to a company's on-premises data center over a Site-to-Site VPN connection. The application is accessible from the company network only. This is a temporary non-production application that is used during business hours. The workload is generally low with occasional surges.
The application has an Amazon Aurora MySQL provisioned database cluster on the backend. The VPC has an internet gateway and a NAT gateways attached. The web servers are in private subnets in an Auto Scaling group behind an Elastic Load Balancer. The web servers also upload data to an Amazon S3 bucket through the internet.
A solutions architect needs to reduce operational costs and simplify the architecture. Which strategy should the solutions architect use?

A. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours onl
B. Use 3-year scheduled Reserved Instances for the web server EC2 instance
C. Detach the internet gateway and remove the NAT gateways from the VP
D. Use an Aurora Servertess database and set up a VPC endpoint for the S3 bucket.
E. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours onl
F. Detach the internet gateway and remove the NAT gateways from the VP
G. Use an Aurora Servertess database and set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.
H. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours onl
I. Detach the internet gateway from the VPC, and use an Aurora Servertess databas
J. Set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.
K. Use 3-year scheduled Reserved Instances for the web server Amazon EC2 instance
L. Remove the NAT gateways from the VPC, and set up a VPC endpoint for the S3 bucke
M. Use Amazon
N. CloudWatch and AWS Lambda to stop and start the Aurora DB cluster so it operates during business hours onl
O. Update the network routing and security rules and policies related to the changes.

Answer: B

Explanation:
The application is accessible from the company network only remove NAT and IGW, application - S3 with VPC endpoint. Non-Production application no need to go for Reserved instances
To build site-to-site vpn, you don't need internet gateway. Instead, customer gateway is needed.
https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html#vpn-create-cgw

NEW QUESTION 29
- (Exam Topic 1)
A company needs to store and process image data that will be uploaded from mobile devices using a custom mobile app. Usage peaks between 8 AM and 5 PM on weekdays, with thousands of uploads per minute. The app is rarely used at any other time A user is notified when image processing is complete.
Which combination of actions should a solutions architect take to ensure image processing can scale to handle the load1? (Select THREE.)

A. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon MQ queue.
B. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue.

C. Invoke an AWS Lambda function to perform image processing when a message is available in the queue.
D. Invoke an S3 Batch Operations job to perform image processing when a message is available in the queue.
E. Send a push notification to the mobile app by using Amazon Simple Notification Service (AmazonSNS) when processing is complete.
F. Send a push notification to the mobile app by using Amazon Simple Email Service (Amazon SES) when processing is complete.

Answer: BCE

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/batch-ops-basics.html

NEW QUESTION 32
- (Exam Topic 1)
A company has has many services running in its on-premises data center. The data center is connected to AWS using AWS Direct Connect (DX) and an IPSec VPN. The service data is sensitive and connectivity cannot traverse the internet. The company wants to expand into a new market segment and begin offering its services to other companies that are using AWS.
Which solution will meet these requirements?

A. Create a VPC Endpoint Service that accepts TCP traffic, host it behind a Network Load Balancer, and make the service available over DX.
B. Create a VPC Endpoint Service that accepts HTTP or HTTPS traffic, host it behind an Application Load Balancer, and make the service available over DX.
C. Attach an internet gateway to the VP
D. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.
E. Attach a NAT gateway to the VP
F. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.

Answer: A

NEW QUESTION 37
- (Exam Topic 1)
An education company is running a web application used by college students around the world. The application runs in an Amazon Elastic Container Service {Amazon ECS) cluster in an Auto Scaling group behind an Application Load Balancer (ALB). A system administrator detects a weekly spike in the number of failed login attempts, which overwhelm the application's authentication service. All the failed login attempts originate from about 500 different IP addresses that change each week, A solutions architect must prevent the failed login attempts from overwhelming the authentication service.
Which solution meets these requirements with the MOST operational efficiency?

A. Use AWS Firewall Manager to create a security group and security group policy to deny access from the IP addresses.
B. Create an AWS WAF web ACL with a rate-based rule, and set the rule action to Bloc
C. Connect theweb ACL to the ALB.
D. Use AWS Firewall Manager to create a security group and security group policy to allow access only to specific CIOR ranges.
E. Create an AWS WAF web ACL with an IP set match rule, and set the rule action to Bloc
F. Connect the web ACL to the ALB.

Answer: B

Explanation:
https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html
The IP set match statement inspects the IP address of a web request against a set of IP addresses and address ranges. Use this to allow or block web requests based on the IP addresses that the requests originate from. By default, AWS WAF uses the IP address from the web request origin, but you can configure the rule to use an HTTP header like X-Forwarded-For instead.
https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-ipset-match.html
https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html

NEW QUESTION 42
- (Exam Topic 1)
A media company uses Amazon DynamoDB to store metadata for its catalog of movies that are available to stream. Each media item Contains user-facing content that concludes a description of the media, a list of search tags, and similar data. In addition, media items include a list of Amazon S3 key names that relate to movie files. The company stores these movie files in a single S3 bucket that has versioning enable. The company uses Amazon CloudFront to serve these movie files.
The company has 100.000 media items, and each media item can have many different S3 objects that represent different encodings of the same media S3 objects that belong to the same media item are grouped together under the same key prefix, which is a random unique ID
Because of an expiring contract with a media provider, the company must remove 2.000 media Items. The company must completely delete all DynamoDB keys and movie files on Amazon S3 that are related to these media items within 36 hours The company must ensure that the content cannot be recovered.
Which combination of actions will meet these requirements? (Select TWO.)

A. Configure the dynamoDB table with a TTL fiel
B. Create and invoke an AWS Lambda function to perform a conditional update Set the TTL field to the time of the contract's expiration on every affected media item.
C. Configure an S3 Lifecycle object expiration rule that is based on the contract's expiration date
D. Write a script to perform a conditional delete on all the affected DynamoDB records
E. Temporarily suspend versioning on the S3 bucke
F. Create and invoke an AWS Lambda function that deletes affected objects Reactivate versioning when the operation is complete
G. Write a script to delete objects from Amazon S3 Specify in each request a NoncurrentVersionExpiration property with a NoncurrentDays attribute set to 0.

Answer: CE

NEW QUESTION 46
- (Exam Topic 1)
A company has an application that generates reports and stores them in an Amazon S3 bucket. When a user accesses their report, the application generates a signed URL to allow the user to download the report. The company's security team has discovered that the files are public and that anyone can download them without authentication. The company has suspended the generation of new reports until the problem is resolved.

Which set of actions will immediately remediate the security issue without impacting the application's normal workflow?

A. Create an AWS Lambda function that applies a deny all policy for users who are not authenticated.Create a scheduled event to invoke the Lambda function.
B. Review the AWS Trusted Advisor bucket permissions check and implement the recommended actions.
C. Run a script that puts a private ACL on all of the objects in the bucket.
D. Use the Block Public Access feature in Amazon S3 to set the IgnorePublicAcls option to TRUE on the bucket.

Answer: D

Explanation:
The S3 bucket is allowing public access and this must be immediately disabled. Setting the IgnorePublicAcls option to TRUE causes Amazon S3 to ignore all public ACLs on a bucket and any objects that it contains. The other settings you can configure with the Block Public Access Feature are:
o BlockPublicAcls – PUT bucket ACL and PUT objects requests are blocked if granting public access. o BlockPublicPolicy – Rejects requests to PUT a bucket policy if granting public access.
o RestrictPublicBuckets – Restricts access to principles in the bucket owners' AWS account. https://aws.amazon.com/s3/features/block-public-access/

NEW QUESTION 51
- (Exam Topic 1)
A large company in Europe plans to migrate its applications to the AWS Cloud. The company uses multiple AWS accounts for various business groups. A data privacy law requires the company to restrict developers'
access to AWS European Regions only.
What should the solutions architect do to meet this requirement with the LEAST amount of management overhead^

A. Create IAM users and IAM groups in each accoun
B. Create IAM policies to limit access to non-European Regions Attach the IAM policies to the IAM groups
C. Enable AWS Organizations, attach the AWS accounts, and create OUs for European Regions andnon-European Region
D. Create SCPs to limit access to non-European Regions and attach the policies to the OUs.
E. Set up AWS Single Sign-On and attach AWS account
F. Create permission sets with policies to restrict access to non-European Regions Create IAM users and IAM groups in each account.
G. Enable AWS Organizations, attach the AWS accounts, and create OUs for European Regions andnon-European Region
H. Create permission sets with policies to restrict access to non-European Region
I. Create IAM users and IAM groups in the primary account.

Answer: B

Explanation:
"This policy uses the Deny effect to deny access to all requests for operations that don't target one of the two approved regions (eu-central-1 and eu-west-1)."
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.htm
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements_condition.html

NEW QUESTION 54
- (Exam Topic 1)
A solutions architect is responsible (or redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.
Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.
Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.
Which approach would improve the availability and durability of (he system while decreasing the processing latency and minimizing costs?

A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda functio
B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SOS queu
D. Extract the core processing code from the existing application and update it to pull items from Amazon SOS instead of an in-memory queu
E. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SOS queue.
F. Modify the application to use Amazon DynamoDB instead of Amazon RD
G. Configure Auto Scaling for the DynamoDB tabl
H. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilizatio
I. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
J. Update the application to use a Redis task queue instead of the in-memory queu
K. 8uild a Docker container image for the applicatio
L. Create an Amazon ECS task definition that includes the application container and a separate container to host Redi
M. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer: B

Explanation:
The obvious challenges here are long workloads, scalability based on queue load, and reliability. Almost always the defacto answer to queue related workload is SQS. Since the workloads are very long (90 minutes) Lambdas cannot be used (15 mins max timeout). So, autoscaled smaller EC2 nodes that wait on external services to complete the task makes more sense. If the task fails, the message is returned to the queue and retried.

NEW QUESTION 56
- (Exam Topic 1)
A company has developed a single-page web application in JavaScript. The source code is stored in a single Amazon S3 bucket in the us-east-1 Region. The company serves the web application to a global user base through Amazon CloudFront.

The company wants to experiment with two versions of the website without informing application users. Each version of the website will reside in its own S3 bucket. The company wants to determine which version is most successful in marketing a new product.

The solution must send application users that are based in Europe to the new website design. The solution must send application users that are based in the United States to the current website design. However, some exceptions exist. The company needs to be able to redirect specific users to the new website design, regardless of the users' location.

Which solution meets these requirements?

A. Configure two CloudFront distribution
B. Configure a geolocation routing policy in Amazon Route 53 to route traffic to the appropriate CloudFront endpoint based on the location of clients.
C. Configure a single CloudFront distributio
D. Create a behavior with different paths for each version of the sit
E. Configure Lambda@Edge on the default path to generate redirects and send the client to the correct version of the website.
F. Configure a single CloudFront distributio
G. Configure an alternate domain name on the distribution.Configure two behaviors to route users to the different S3 origins based on the domain name that the client uses in the HTTP request.
H. Configure a single CloudFront distribution with Lambda@Edg
I. Use Lambda@Edge to send user requests to different origins based on request attributes.

Answer: A

NEW QUESTION 58
- (Exam Topic 1)
A company is building an image service on the web that will allow users to upload and search random photos. At peak usage, up to 10.000 users worldwide will upload their images. The service will then overlay text on the uploaded images, which will then be published on the company website.
Which design should a solutions architect implement?

A. Store the uploaded images in Amazon Elastic File System (Amazon EFS). Send application log information about each image to Amazon CloudWatch Log
B. Create a fleet of Amazon EC2 instances that use CloudWatch Logs to determine which images need to be processe
C. Place processed images in anolher directory in Amazon EF
D. Enable Amazon CloudFront and configure the origin to be the one of the EC2 instances in the fleet.
E. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to Amazon Simple Notification Service (Amazon SNS). Create a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) to pull messages from Amazon SNS to process the images and place them in Amazon Elastic File System (Amazon EFS). Use Amazon CloudWatch metrics for the SNS message volume to scale out EC2 instance
F. Enable Amazon CloudFront and configure the origin lo be the ALB in front of the EC2 instances.
G. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to the Amazon Simple Queue Service (Amazon SOS) queu
H. Create a fleet of Amazon EC2 instances to pull messages from Ihe SOS queue to process the images and place them in another S3 bucke
I. Use Amazon CloudWatch metrics for queue depth to scale out EC2 instance
J. Enable Amazon CloudFront and configure the origin to be the S3 bucket that contains the processed images.
K. Store the uploaded images on a shared Amazon Elastic Block Store (Amazon EBS) volume mounted toa fleet of Amazon EC2 Spot instance
L. Create an Amazon DynamoDB table that contains information about each uploaded image and whether it has been processe
M. Use an Amazon EventBridge (Amazon CloudWatch Events) rule lo scale out EC2 instance
N. Enable Amazon CloudFront and configure the origin to reference an Elastic Load Balancer in front of the fleet of EC2 instances.

Answer: C

NEW QUESTION 60
- (Exam Topic 1)
A company plans to migrate to AWS. A solutions architect uses AWS Application Discovery Service over the fleet and discovers that there is an Oracle data warehouse and several PostgreSQL databases. Which combination of migration patterns will reduce licensing costs and operational overhead? (Select TWO.)

A. Lift and shift the Oracle data warehouse to Amazon EC2 using AWS DMS.
B. Migrate the Oracle data warehouse to Amazon Redshift using AWS SCT and AWS QMS.
C. Lift and shift the PostgreSQL databases to Amazon EC2 using AWS DMS.
D. Migrate the PostgreSQL databases to Amazon RDS for PostgreSQL using AWS DMS
E. Migrate the Oracle data warehouse to an Amazon EMR managed cluster using AWS DMS.

Answer: BD

Explanation:
https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/ https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-an-on-premises-postgresql-database

NEW QUESTION 64
- (Exam Topic 1)
A team collects and routes behavioral data for an entire company. The company runs a Multi-AZ VPC environment with public subnets, private subnets, and in internet gateway Each public subnet also contains a NAT gateway Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads run in private subnets.
A solutions architect must review the infrastructure The solutions architect needs to reduce costs and maintain the function of the applications. The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.
What should the solutions architect do to meet these requirements?

A. Enable VPC Flow Log
B. Use Amazon Athena to analyze the logs for traffic that can be remove
C. Ensure that security groups are blocking traffic that is responsible for high costs.
D. Add an interface VPC endpoint for Kinesis Data Streams to the VP
E. Ensure that applications have thecorrect IAM permissions to use the interface VPC endpoint.
F. Enable VPC Flow Logs and Amazon Detectiv
G. Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic

H. Add an interface VPC endpoint for Kinesis Data Streams to the VPC Ensure that the VPC endpoint policy allows traffic from the applications

Answer: D

Explanation:

https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/
VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint

NEW QUESTION 66
- (Exam Topic 1)
A solutions architect is evaluating the reliability of a recently migrated application running on AWS. The front end is hosted on Amazon S3 and accelerated by Amazon CloudFront. The application layer is running in a stateless Docker container on an Amazon EC2 On-Demand Instance with an Elastic IP address. The storage layer is a MongoDB database running on an EC2 Reserved Instance in the same Availability Zone as the application layer.
Which combination of steps should the solutions architect take to eliminate single points of failure with minimal application code changes? (Select TWO.)

A. Create a REST API in Amazon API Gateway and use AWS Lambda functions as the application layer.
B. Create an Application Load Balancer and migrate the Docker container to AWS Fargate.
C. Migrate the storage layer to Amazon DynamoD8.
D. Migrate the storage layer to Amazon DocumentD8 (with MongoDB compatibility).
E. Create an Application Load Balancer and move the storage layer to an EC2 Auto Scaling group.

Answer: BD

Explanation:

https://aws.amazon.com/documentdb/?nc1=h_ls
https://aws.amazon.com/blogs/containers/using-alb-ingress-controller-with-amazon-eks-on-fargate/

NEW QUESTION 68
- (Exam Topic 1)
A company runs an application that gives users the ability to search for videos and related information by using keywords that are curated from content providers. The application data is stored in an on-premises Oracle database that is 800 GB in size.
The company wants to migrate the data to an Amazon Aurora MySQL DB instance. A solutions architect
plans to use the AWS Schema Conversion Tool and AWS Database Migration Service (AWS DMS) for the migration. During the migration, the existing database must serve ongoing requests. The migration must be completed with minimum downtime
Which solution will meet these requirements?

A. Create primary key indexes, secondary indexes, and referential integrity constraints in the target database before starting the migration process
B. Use AWS DMS to run the conversion report for Oracle to Aurora MySQ
C. Remediate any issues Then use AWS DMS to migrate the data
D. Use the M5 or CS DMS replication instance type for ongoing replication
E. Turn off automatic backups and logging of the target database until the migration and cutover processes are complete

Answer: B

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html

NEW QUESTION 71
- (Exam Topic 1)
A company provides a centralized Amazon EC2 application hosted in a single shared VPC. The centralized application must be accessible from client applications running in the VPCs of other business units. The centralized application front end is configured with a Network Load Balancer (NLB) for scalability.
Up to 10 business unit VPCs will need to be connected to the shared VPC. Some of the business unit VPC CIDR blocks overlap with the shared VPC. and some overlap with each other. Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only.
Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

A. Create an AW5 Transit Gatewa
B. Attach the shared VPC and the authorized business unit VPCs to the transit gatewa
C. Create a single transit gateway route table and associate it with all of the attached VPC
D. Allow automatic propagation of routes from the attachments into the route tabl
E. Configure VPC routing tables to send traffic to the transit gateway.
F. Create a VPC endpoint service using the centralized application NLB and enable (he option to require endpoint acceptanc
G. Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint servic
H. Accept authorized endpoint requests from the endpoint service console.
I. Create a VPC peering connection from each business unit VPC to Ihe shared VP
J. Accept the VPC peering connections from the shared VPC consol
K. Configure VPC routing tables to send traffic to the VPC peering connection.
L. Configure a virtual private gateway for the shared VPC and create customer gateways for each of theauthorized business unit VPC
M. Establish a Sile-to-Site VPN connection from the business unit VPCs to the shared VP
N. Configure VPC routing tables to send traffic to the VPN connection.

Answer: B

Explanation:

Amazon Transit Gateway doesn't support routing between Amazon VPCs with overlapping CIDRs. If you attach a new Amazon VPC that has a CIDR which overlaps with an already attached Amazon VPC, Amazon Transit Gateway will not propagate the new Amazon VPC route into the Amazon Transit Gateway route table.
https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#client-ip-pre

NEW QUESTION 76
- (Exam Topic 1)
A solutions architect works for a government agency that has strict disaster recovery requirements All Amazon Elastic Block Store (Amazon EBS) snapshots are required to be saved in at least two additional AWS Regions. The agency also is required to maintain the lowest possible operational overhead.
Which solution meets these requirements?

A. Configure a policy in Amazon Data Lifecycle Manager (Amazon DLMJ to run once daily to copy the EBS snapshots to the additional Regions.
B. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to copy the EBS snapshots to the additional Regions.
C. Set up AWS Backup to create the EBS snapshot
D. Configure Amazon S3 cross-Region replication to copy the EBS snapshots to the additional Regions.
E. Schedule Amazon EC2 Image Builder to run once daily to create an AMI and copy the AMI to the additional Regions.

Answer: B

NEW QUESTION 78
- (Exam Topic 1)
A financial company is building a system to generate monthly, immutable bank account statements for its users. Statements are stored in Amazon S3. Users should have immediate access to their monthly statements for up to 2 years. Some users access their statements frequently, whereas others rarely access their statements. The company's security and compliance policy requires that the statements be retained for at least 7 years.
What is the MOST cost-effective solution to meet the company's needs?

A. Create an S3 bucket with Object Lock disable
B. Store statements in S3 Standar
C. Define an S3 Lifecycle policy to transition the data to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 day
D. Define another S3 Lifecycle policy to move the data to S3 Glacier Deep Archive after 2 year
E. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
F. Create an S3 bucket with versioning enable
G. Store statements in S3 Intelligent-Tierin
H. Usesame-Region replication to replicate objects to a backup S3 bucke
I. Define an S3 Lifecycle policy for the backup S3 bucket to move the data to S3 Glacie
J. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
K. Create an S3 bucket with Object Lock enable
L. Store statements in S3 Intelligent-Tierin
M. Enable compliance mode with a default retention period of 2 year
N. Define an S3 Lifecycle policy to move the data to S3 Glacier after 2 year
O. Attach an S3 Glacier Vault Lock policy with deny delete permissionsfor archives less than 7 years old.
P. Create an S3 bucket with versioning disable
Q. Store statements in S3 One Zone-Infrequent Access (S3 One Zone-IA). Define an S3 Lifecyde policy to move the data to S3 Glacier Deep Archive after 2 year
R. Attach an S3 Glader Vault Lock policy with deny delete permissions for archives less than 7 years old.

Answer: C

Explanation:
https://aws.amazon.com/about-aws/whats-new/2018/11/s3-object-lock/
Create an S3 bucket with Object Lock enabled. Store statements in S3 Intelligent-Tiering. Enable compliance mode with a default retention period of 2 years.
Define an S3 Lifecycle policy to move the data to S3 Glacier after 2 years. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html

NEW QUESTION 83
- (Exam Topic 1)
A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.
After initial deployment, the company observes 1.000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.
Which approach should the company take to secure its API?

A. Create an Amazon CloudFront distribution with the API as the origi
B. Create an AWS WAF web ACL with a rule to block clients "hat submit more than five requests per da
C. Associate the web ACL with the CloudFront distributio
D. Configure CloudFront with an origin access identity (OAI) and associate it with the distributio
E. Configure API Gateway to ensure only the OAI can execute the POST method.
F. Create an Amazon CloudFront distribution with the API as the origi
G. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per da
H. Associate the web ACL with the CloudFront distributio
I. Add a custom header to the CloudFront distribution populated with an API ke
J. Configure the API to require an API key on the POST method.
K. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners.Associate the web ACL with the AP
L. Create a resource policy with a request limit and associate it with the AP
M. Configure the API to require an API key on the POST method.
N. Associate the web ACL with the AP
O. Create a usage plan with a request limit and associate it with the AP
P. Create an API key and add it to the usage plan.

Answer: D

Explanation:
"A usage plan specifies who can access one or more deployed API stages and methods—and also how much and how fast they can access them. The plan uses API keys to identify API clients and meters access to the associated API stages for each key. It also lets you configure throttling limits and quota limits that are enforced on individual client API keys."

https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html

NEW QUESTION 88
- (Exam Topic 1)
A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.
What should a solutions architect do to meet these requirements?

A. Create a new developer accoun
B. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organization
C. Enforce a tagging policy that denotes Region affinity.
D. Create an SCP that denies the launch of all EC2 instances except I3.small EC2 instances in us-east-2.Attach the SCP to the project's account.
E. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.
F. Create an IAM policy than allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

Answer: D

NEW QUESTION 89
- (Exam Topic 1)
A company is serving files to Its customers through an SFTP server that is accessible over the internet The SFTP server is running on a single Amazon EC2 instance with an Elastic IP address attached Customers connect to the SFTP server through its Elastic IP address and use SSH (or authentication. The EC2 instance also has an attached security group that allows access from all customer IP addresses.
A solutions architect must implement a solution to improve availability, minimize the complexity of infrastructure management, and minimize the disruption to customers who access files The solution must not change the way customers connect.
Which solution will meet these requirements?

A. Disassociate the Elastic IP address from the EC2 instanc
B. Create an Amazon S3 bucket to be used for SFTP file hostin
C. Create an AWS Transfer Family server Configure the Transfer Family server with a publicly accessible endpoint Associate the SFTP Elastic IP address with the new endpoint Point the Transfer Family server to the S3 bucke
D. Sync all files from the SFTP server to the S3 bucket.
E. Disassociate the Elastic IP address from the EC2 instanc
F. Create an Amazon S3 bucket to be used for SFTP file hostin
G. Create an AWS Transfer Family serve
H. Configure the Transfer Family server with aVPC-hoste
I. internet-facing endpoin
J. Associate the SFTP Elastic IP address with the new endpoin
K. Attach the security group with customer IP addresses to the new endpoin
L. Point the Transfer Family server to the S3 bucket Sync all files from the SFTP server to the S3 bucket.
M. Disassociate the Elastic IP address from the EC2 instanc
N. Create a new Amazon Elastic File System{Amazon EFS) file system to be used for SFTP file hostin
O. Create an AWS Fargate task definition to run an SFTP serve
P. Specify the EFS file system as a mount in the task definitio
Q. Create a Fargate service by using the task definition, and place a Network Load Balancer (NLB) in front of the service When configuring the service, attach the security group with customer IP addresses to the tasks that run the SFTP serve
R. Associate the Elastic IP address with the NL
S. Sync all files from the SFTP server to the S3 bucket.
T. Disassociate the Elastic IP address from the EC2 instanc
. Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hostin
. Create a Network Load Balancer (NLB) with the Elastic IP address attache
. Create an Auto Scaling group with EC2 instances that run an SFTP server Define in the Auto Scaling group that instances that are launched should attach the newmulti-attach EBS volume Configure the Auto Scaling group to automatically add instances behind the NLB Configure the Auto Scaling group to use the security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launche
. Sync all files from the SFTP server to the new multi-attach EBS volume.

Answer: B

Explanation:
https://docs.aws.amazon.com/transfer/latest/userguide/create-server-in-vpc.html https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/

NEW QUESTION 93
- (Exam Topic 1)
A solution architect needs to deploy an application on a fleet of Amazon EC2 instances. The EC2 instances run in private subnets in An Auto Scaling group. The application is expected to generate logs at a rate of 100 MB each second on each of the EC2 instances.
The logs must be stored in an Amazon S3 bucket so that an Amazon EMR cluster can consume them for further processing The logs must be quickly accessible for the first 90 days and should be retrievable within 48 hours thereafter.
What is the MOST cost-effective solution that meets these requirements?

A. Set up an S3 copy job to write logs from each EC2 instance to the S3 bucket with S3 Standard storage Use a NAT instance within the private subnets to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier.
B. Set up an S3 sync job to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage Use a gateway VPC endpoint for Amazon S3 to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier Deep Archive
C. Set up an S3 batch operation to copy logs from each EC2 instance to the S3 bucket with S3 Standardstorage Use a NAT gateway with the private subnets to connect to Amazon S3 Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier Deep Archive
D. Set up an S3 sync job to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage Use a gateway VPC endpoint for Amazon S3 to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier

Answer: C

NEW QUESTION 95
- (Exam Topic 1)
A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance The DB instance is expected to receive many more reads than writes. The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available.
Which steps should the solutions architect take to meet these requirements? (Select THREE)

A. Create multiple read replicas and put them into an Auto Scaling group.
B. Create multiple read replicas in different Availability Zones.
C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy.
D. Create an Application Load Balancer (ALB) and put the read replicas behind the ALB.
E. Configure an Amazon CloudWatch alarm to detect a failed read replic
F. Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.
G. Configure an Amazon Route 53 health check for each read replica using its endpoint

Answer: BCF

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/requests-rds-read-replicas/
You can use Amazon Route 53 weighted record sets to distribute requests across your read replicas. Within a Route 53 hosted zone, create individual record sets for each DNS endpoint associated with your read replicas and give them the same weight. Then, direct requests to the endpoint of the record set. You can incorporate Route 53 health checks to be sure that Route 53 directs traffic away from unavailable read replicas

NEW QUESTION 96
- (Exam Topic 1)
A medical company is running a REST API on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group behind an Application Load Balancer (ALB). The ALB runs in three public subnets, and the EC2 instances run in three private subnets. The company has deployed an Amazon CloudFront distribution that has the AL8 as the only origin.
Which solution should a solutions architect recommend to enhance the origin security?

A. Store a random string in AWS Secrets Manage
B. Create an AWS Lambda (unction for automatic secret rotatio
C. Configure CloudFront to inject the random string as a custom HTTP header for the origin reques
D. Create an AWS WAF web ACL rule with a string match rule for the custom heade
E. Associate the web ACL with the ALB.
F. Create an AWS WAF web ACL rule with an IP match condition of the CloudFront service IP address range
G. Associate the web ACL with the AL
H. Move the ALB into the three private subnets.
I. Store a random string in AWS Systems Manager Parameter Stor
J. Configure Parameter Store automatic rotation for the strin
K. Configure CloudFront to inject the random siring as a custom HTTP header for the origin reques
L. Inspect the value of the custom HTTP header, and block access in the ALB.
M. Configure AWS Shield Advance
N. Create a security group policy to allow connections from CloudFront service IP address range
O. Add the policy to AWS Shield Advanced, and attach the policy to the ALB.

Answer: D

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html
it shows For Amazon EC2 Auto Scaling, there are two primary process types: Launch and Terminate. The Launch process adds a new Amazon EC2 instance to an Auto Scaling group, increasing its capacity. The Terminate process removes an Amazon EC2 instance from the group, decreasing its capacity. HealthCheck process for EC2 autoscaling is not a primary process! It is a process along with the following AddToLoadBalancer AlarmNotification AZRebalance HealthCheck InstanceRefresh ReplaceUnhealthy ScheduledActions From the requirements, Some EC2 instances are now being marked as unhealthy and are being terminated. Application is running at reduced capacity not because instances are marked unhealthy but because they are being terminated.
https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html#choosing-suspend-r

NEW QUESTION 99
- (Exam Topic 1)
An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list.
The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets.
How should a solutions architect ensure that the web application can continue to call the third-parly API after the migration?

A. Associate a block of customer-owned public IP addresses to the VP
B. Enable public IP addressing for public subnets in the VPC.
C. Register a block of customer-owned public IP addresses in the AWS accoun
D. Create Elastic IP addresses from the address block and assign them lo the NAT gateways in the VPC.
E. Create Elastic IP addresses from the block of customer-owned IP addresse
F. Assign the static Elastic IP addresses to the ALB.
G. Register a block of customer-owned public IP addresses in the AWS accoun
H. Set up AWS Global Accelerator to use Elastic IP addresses from the address bloc
I. Set the ALB as the accelerator endpoint.

Answer: B

Explanation:

When EC2 instances reach third-party API through internet, their privates IP addresses will be masked by NAT Gateway public IP address.
https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-bring-your-own-ip-byoip-for-amaz

NEW QUESTION 103
- (Exam Topic 1)
A company needs to run a software package that has a license that must be run on the same physical host for the duration of Its use. The software package is only going to be used for 90 days The company requires patching and restarting of all instances every 30 days
How can these requirements be met using AWS?

A. Run a dedicated instance with auto-placement disabled.
B. Run the instance on a dedicated host with Host Affinity set to Host.
C. Run an On-Demand Instance with a Reserved Instance to ensure consistent placement.
D. Run the instance on a licensed host with termination set for 90 days.

Answer: B

Explanation:
Host Affinity is configured at the instance level. It establishes a launch relationship between an instance and a Dedicated Host. (This set which host the instance can run on) Auto-placement allows you to manage whether instances that you launch are launched onto a specific host, or onto any available host that has matching configurations. Auto-placement must be configured at the host level. (This sets which instance the host can run.) When affinity is set to Host, an instance launched onto a specific host always restarts on the same host if stopped. This applies to both targeted and untargeted launches.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/how-dedicated-hosts-work.html
When affinity is set to Off, and you stop and restart the instance, it can be restarted on any available host. However, it tries to launch back onto the last Dedicated Host on which it ran (on a best-effort basis).

NEW QUESTION 106
- (Exam Topic 1)
A company has a data lake in Amazon S3 that needs to be accessed by hundreds of applications across many AWS accounts. The company's information security policy states that the S3 bucket must not be accessed over the public internet and that each application should have the minimum permissions necessary to function.
To meet these requirements, a solutions architect plans to use an S3 access point that is restricted to specific VPCs tor each application.
Which combination of steps should the solutions architect take to implement this solution? (Select TWO.)

A. Create an S3 access point for each application in the AWS account that owns the S3 bucke
B. Configure each access point to be accessible only from the application's VP
C. Update the bucket policy to require access from an access point.
D. Create an interface endpoint for Amazon S3 in each application's VP
E. Configure the endpoint policy to allow access to an S3 access poin
F. Create a VPC gateway attachment for the S3 endpoint.
G. Create a gateway endpoint lor Amazon S3 in each application's VP
H. Configure the endpoint policy to allow access to an S3 access poin
I. Specify the route table that is used to access the access point.
J. Create an S3 access point for each application in each AWS account and attach the access points to the S3 bucke
K. Configure each access point to be accessible only from the application's VP
L. Update the bucket policy to require access from an access point.
M. Create a gateway endpoint for Amazon S3 in the data lake's VP
N. Attach an endpoint policy to allow access to the S3 bucke
O. Specify the route table that is used to access the bucket.

Answer: AC

Explanation:
https://joe.blog.freemansoft.com/2020/04/protect-data-in-cloud-with-s3-access.html https://aws.amazon.com/s3/features/access-points/
https://aws.amazon.com/s3/features/access-points/
&
https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/

NEW QUESTION 110
- (Exam Topic 1)
A developer reports receiving an Error 403: Access Denied message when they try to download an object from an Amazon S3 bucket. The S3 bucket is accessed using an S3 endpoint inside a VPC. and is encrypted with an AWS KMS key. A solutions architect has verified that (he developer is assuming the correct IAM role in the account that allows the object to be downloaded. The S3 bucket policy and the NACL are also valid.
Which additional step should the solutions architect take to troubleshoot this issue?

A. Ensure that blocking all public access has not been enabled in the S3 bucket.
B. Verify that the IAM rote has permission to decrypt the referenced KMS key.
C. Verify that the IAM role has the correct trust relationship configured.
D. Check that local firewall rules are not preventing access to the S3 endpoint.

Answer: B

NEW QUESTION 112
- (Exam Topic 1)
A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalogue page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.
Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times.
Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire

application stack for future growth? (Select TWO.)

A. Configure read replicas for Amazon RDS MySQL and use the single reader endpoint in the web application to reduce the load on the backend database tier.
B. Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionalit
C. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
D. Configure the target group health check to use a TCP check of the Amazon EC2 web server and theAmazon Route 53 health check against the product page to evaluate full application functionalit
E. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
F. Configure an Amazon CloudWatch alarm for Amazon RDS with an action to recover a high-load, impaired RDS instance in the database tier.
G. Configure an Amazon ElastiCache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

Answer: BE

Explanation:
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/health-checks-types.html

NEW QUESTION 115
- (Exam Topic 2)
A finance company is storing financial records in an Amazon S3 bucket. The company persists a record for every financial transaction. According to regulatory requirements, the records cannot be modified for at least 1 year after they are written. The records are read on a regular basis and must be immediately accessible.
Which solution will meet these requirements?

A. Create a new S3 bucke
B. Turn on S3 Object Lock, set a default retention period of 1 year, and set the retention mode to compliance mod
C. Store all records inthe new S3 bucket.
D. Create an S3 Lifecycle rule to immediately transfer new objects to the S3 Glacier storage tier Create an S3 Glacier Vault Lock policy that has a retention period of 1 year.
E. Create an S3 Lifecycle rule to immediately transfer new objects to the S3 Intelligent-Tiering storage tier.Set a retention period of 1 year.
F. Create an S3 bucket policy with a Deny action for PutObject operations with a condition where the s3:x-amz-object-retention header is not equal to 1 year.

Answer: A

NEW QUESTION 117
- (Exam Topic 2)
A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days
The company has a high-speed AWS Direct Connect connection Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day Which solution meets these requirements?

A. Use regularly scheduled AWS Snowball Edge devices to transfer the sequencing data into AWS When AWS receives the Snowball Edge device and the data is loaded into Amazon S3 use S3 events to trigger an AWS Lambda function to process the data
B. Use AWS Data Pipeline to transfer the sequencing data to Amazon S3 Use S3 events to trigger an Amazon EC2 Auto Scaling group to launch custom-AMI EC2 instances running the Docker containers to process the data
C. Use AWS DataSync to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Lambda function that starts an AWS Step Functions workflow Store the Docker images in Amazon Elastic Container Registry (Amazon ECR) and trigger AWS Batch to run the container and process the sequencing data
D. Use an AWS Storage Gateway file gateway to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Batch job that runs on Amazon EC2 instances running the Docker containers to process the data

Answer: C

NEW QUESTION 118
- (Exam Topic 2)
A company hosts a blog post application on AWS using Amazon API Gateway. Amazon DynamoDB, and AWS Lambda The application currently does not use API keys to authorize requests The API model is as follows:
GET /posts/Jpostld) to get post details
GET /users/{userld}. to get user details
GET /comments/{commentld}: to get comments details
The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by making the comments appear in real time
Which design should be used to reduce comment latency and improve user experience?

A. Use edge-optimized API with Amazon CloudFront to cache API responses.
B. Modify the blog application code to request GET/commentsV{commentld} every 10 seconds
C. Use AWS AppSync and leverage WebSockets to deliver comments
D. Change the concurrency limit of the Lambda functions to lower the API response time.

Answer: C

NEW QUESTION 122
- (Exam Topic 2)
A company that develops consumer electronics with offices in Europe and Asia has 60 TB of software images stored on premises in Europe. The company wants to transfer the images to an Amazon S3 bucket in the ap-northeast-1 Region. New software images are created daily and must be encrypted in transit. The

company needs a solution that does not require custom development to automatically transfer all existing and new software images to Amazon S3.
What is the next step in the transfer process?

A. Deploy an AWS DataSync agent and configure a task to transfer the images to the S3 bucket.
B. Configure Amazon Kinesis Data Firehose to transfer the images using S3 Transfer Acceleration.
C. Use an AWS Snowball device to transfer the images with the S3 bucket as the target.
D. Transfer the images over a Site-to-Site VPN connection using the S3 API with multipart upload.

Answer: A

NEW QUESTION 126
- (Exam Topic 2)
A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet 172.31.0.0/24 within VPC A with CIDR block 172.31.0.0/16. The developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block 172.50.0.0/16. The IP of the database instance is hard-coded in the application instances.
Which recommendations should a solutions architect present to the developers to solve the problem in a secure way with minimal maintenance and overhead''

A. Disable the SrcDestCheck attribute for all instances running the application and Oracle Database.Change the default route of VPC A to point ENI of the Oracle Database that has an IP address assigned within the range of 172.50.0.0/16
B. Create and attach internet gateways for both VPC
C. Configure default routes to the internet gateways for both VPC
D. Assign an Elastic IP for each Amazon EC2 instance in VPC A
E. Create a VPC peering connection between the two VPCs and add a route to the routing table of VPC A that points to the IP address range of 172.50.0.0/16
F. Create an additional Amazon EC2 instance for each VPC as a customer gateway; create one virtual private gateway (VGW) for each VP
G. configure an end-to-end VPC, and advertise the routes for 172.50.0.0/16

Answer: C

NEW QUESTION 128
- (Exam Topic 2)
A company that runs applications on AWS recently subscribed to a new software-as-a-service (SaaS) data vendor. The vendor provides the data by way of a REST API that the vendor hosts in its AWS environment The vendor offers multiple options for connectivity to the API and Is working with the company to find the best way to connect.
The company's AWS account does not allow outbound internet access from Its AWS environment The vendor's services run on AWS in the same AWS Region as the company's applications
A solutions architect must Implement connectivity to the vendor's API so that the API is highly available In the company's VPC.
Which solution will meet these requirements?

A. Connect to the vendor's public API address for the data service.
B. Connect to the vendor by way of a VPC peering connection between the vendor's VPC and the company's VPC
C. Connect to the vendor by way of a VPC endpoint service that uses AWS PrivateLink
D. Connect to a public bastion host that the vendor provides Tunnel the API traffic.

Answer: C

NEW QUESTION 133
- (Exam Topic 2)
A company is running a three-tier web application in an on-premises data center. The frontend is served by an Apache web server, the middle tier is a monolithic Java application, and the storage tier is a PostgreSOL database.
During a recent marketing promotion, customers could not place orders through the application because the application crashed An analysis showed that all three tiers were overloaded. The application became unresponsive, and the database reached its capacity limit because of read operations. The company already has several similar promotions scheduled in the near future.
A solutions architect must develop a plan for migration to AWS to resolve these issues. The solution must maximize scalability and must minimize operational effort.
Which combination of steps will meet these requirements? (Select THREE.)

A. Refactor the frontend so that static assets can be hosted on Amazon S3. Use Amazon CloudFront to serve the frontend to customer
B. Connect the frontend to the Java application.
C. Rehost the Apache web server of the frontend on Amazon EC2 instances that are in an Auto Scaling grou
D. Use a load balancer in front of the Auto Scaling grou
E. Use Amazon Elastic File System (Amazon EFS) to host the static assets that the Apache web server needs.
F. Rehost the Java application in an AWS Elastic Beanstalk environment that includes auto scaling.
G. Refactor the Java applicatio
H. Develop a Docker container to run the Java applicatio
I. Use AWS Fargate to host the container.
J. Use AWS Database Migration Service (AWS DMS) to replatform the PostgreSQL database to an Amazon Aurora PostgreSQL databas
K. Use Aurora Auto Scaling for read replicas.
L. Rehost the PostgreSQL database on an Amazon EC2 instance that has twice as much memory as the on-premises server.

Answer: BCF

NEW QUESTION 135
- (Exam Topic 2)
A development team s Deploying new APIs as serverless applications within a company. The team is currently using the AWS Management Console to provision Amazon API Gateway. AWS Lambda, and Amazon DynamoDB resources A solutions architect has been tasked with automating the future deployments of these serveriess APIs
How can this be accomplished?

A. Use AWS CloudFonTiation with a Lambda-backed custom resource to provision API Gateway Use the MfS: :OynMoDB::Table and AWS::Lambda::Function resources to create the Amazon DynamoOB table and Lambda functions Write a script to automata the deployment of the CloudFormation template.
B. Use the AWS Serverless Application Model to define the resources Upload a YAML template and application files to the code repository Use AWS CodePipeline to conned to the code repository and to create an action to build using AWS CodeBuil
C. Use the AWS CloudFormabon deployment provider m CodePipeline to deploy the solution.
D. Use AWS CloudFormation to define the serverless applicatio
E. Implement versioning on the Lambda functions and create aliases to point to the version
F. When deploying, configure weights to implement shifting traffic to the newest version, and gradually update the weights as traffic moves over
G. Commit the application code to the AWS CodeCommit code repositor
H. Use AWS CodePipeline and connect to the CodeCommit code repository Use AWS CodeBuild to build and deploy the Lambda functions using AWS CodeDeptoy Specify the deployment preference type in CodeDeploy to gradually shift traffic over to the new version.

Answer: B

NEW QUESTION 136
- (Exam Topic 2)
A company is planning to migrate an application from on premises to the AWS Cloud. The company will begin the migration by moving the application's underlying data storage to AWS The application data is stored on a shared tie system on premises, and the application servers connect to the shared We system through SMB.
A solutions architect must implement a solution that uses an Amazon S3 bucket tor shared storage Until the application Is fully migrated and code is rewritten to use native Amazon S3 APIs, the application must continue to have access to the data through SMB The solutions architect must migrate the application data to AWS to its new location while still allowing the on-premises application to access the data.
Which solution will meet these requirements?

A. Create a new Amazon FSx for Windows File Server fie system Configure AWS DataSync with one location tor the on-premises file share and one location for the new Amazon FSx file system Create a new DataSync task to copy the data from the on-premises file share location to the Amazon FSx file system
B. Create an S3 bucket for the applicatio
C. Copy the data from the on-premises storage to the S3 bucket
D. Deploy an AWS Server Migration Service (AWS SMS) VM to the on-premises environmen
E. Use AWS SMS to migrate the file storage server from on premises to an Amazon EC2 instance
F. Create an S3 bucket for the applicatio
G. Deploy a new AWS Storage Gateway Me gateway on anon-premises V
H. Create a new file share that stores data in the S3 bucket and is associated with the tie gatewa
I. Copy the data from the on-premises storage to the new file gateway endpoint.

Answer: A

NEW QUESTION 139
- (Exam Topic 2)
A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts AWS Site-to-Site VPN connections are configured between ail of the company's global offices and the transit account The company has AWS Config enabled on all of its accounts.
The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices Developers Will reference this list to gain access to applications securely.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is update
B. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with Vie updated IP address ranges.
C. Create a new AWS Config managed rule that contains all of the internal IP address ranges Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address range
D. Configure the rule to automatically remediate any noncompliant security group that is detected.
E. In the transit account, create a VPC prefix list with all of the internal IP address range
F. Use AWS Resource Access Manager to share the prefix list with all of the other account
G. Use the shared prefix list to configure security group rules is the other accounts.
H. In the transit account create a security group with all of the internal IP address range
I. Configure the security groups in me other accounts to reference the transit account's securitygroup by using a nested security group reference of *<transit-account-id>./sg-1a2b3c4d".

Answer: C

NEW QUESTION 141
- (Exam Topic 2)
A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.
How should a solutions architect configure the web ACLs to meet these requirements?

A. Set the action of the web ACL rules to Coun
B. Enable AWS WAF logging Analyze the requests for false positives Modify the rules to avoid any false positive Over time change the action of the web ACL rules from Count to Block.
C. Use only rate-based rules in the web ACL
D. and set the throttle limit as high as possible Temporarilyblock all requests that exceed the limi
E. Define nested rules to narrow the scope of the rate tracking.
F. Set the action o' the web ACL rules to Bloc
G. Use only AWS managed rule groups in the web ACLs Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.
H. Use only custom rule groups in the web ACL
I. and set the action to Allow Enable AWS WAF logging Analyze the requests tor false positives Modify the rules to avoid any false positive Over time, change the action of the web ACL rules from Allow to Block.

Answer: B

NEW QUESTION 145
- (Exam Topic 2)
A company has deployed an application to multiple environments in AWS. including production and testing the company has separate accounts for production and testing, and users are allowed to create additional
application users for team members or services. as needed. The security team has asked the operations team tor better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments
Which of the following options would MOST securely accomplish this goal?

A. Create a new AWS account to hold user and service accounts, such as an identity account Create users and groups m the identity accoun
B. Create roles with appropriate permissions in the production and testing accounts Add the identity account to the trust policies for the roles
C. Modify permissions in the production and testing accounts to limit creating new IAM users to members of the operations team Set a strong IAM password policy on each account Create new IAM users and groups in each account to Limit developer access to just the services required to complete their job function.
D. Create a script that runs on each account that checks user accounts For adherence to a security policy.Disable any user or service accounts that do not comply.
E. Create all user accounts in the production account Create roles for access in me production account and testing account
F. Grant cross-account access from the production account to the testing account

Answer: A

NEW QUESTION 147
- (Exam Topic 2)
A retail company is running an application that stores invoice files in an Amazon S3 bucket and metadata about the files in an Amazon DynamoDB table. The application software runs in both us-east-1 and eu-west-1 The S3 bucket and DynamoDB table are in us-east-1. The company wants to protect itself from data corruption and loss of connectivity to either Region
Which option meets these requirements?

A. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Enable versioning on the S3 bucket
B. Create an AWS Lambda function triggered by Amazon CloudWatch Events to make regular backups of the DynamoDB table Set up S3 cross-region replication from us-east-1 to eu-west-1 Set up MFA delete on the S3 bucket in us-east-1.
C. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable versioning on the S3 bucket Implement strict ACLs on the S3 bucket
D. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Set up S3 cross-region replication from us-east-1 toeu-west-1.

Answer: B

NEW QUESTION 149
- (Exam Topic 2)
A company is running an application in the AWS Cloud. The application runs on containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost. Which solution will meet these requirements with the LEAST amount of operational overhead?

A. Provision an Aurora Replica in a different Region.
B. Set up AWS DataSync for continuous replication of the data to a different Region.
C. Set up AWS Database Migration Service (AWS DMS) to perform a continuous replication of the data to a different Region.
D. Use Amazon Data Lifecycle Manager {Amazon DLM) to schedule a snapshot every 5 minutes.

Answer: B

NEW QUESTION 154
- (Exam Topic 2)
A large company recently experienced an unexpected increase in Amazon RDS and Amazon DynamoDB costs The company needs to increase visibility into details of AWS Billing and Cost Management There are various accounts associated with AWS Organizations, including many development and production accounts. There is no consistent tagging strategy across the organization, but there are guidelines in place that require all infrastructure to be deployed using AWS Cloud Formation with consistent tagging Management requires cost center numbers and project ID numbers for all existing and future DynamoDB tables and RDS instances
Which strategy should the solutions architect provide to meet these requirements?

A. Use Tag Editor to tag existing resources Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources
B. Use an AWS Config rule to alert the finance team of untagged resources Create a centralized AWS Lambda based solution to tag untagged RDS databases and DynamoDB resources every hour using a cross-account rote.
C. Use Tag Editor to tag existing resources Create cost allocation tags to define the cost center and project ID Use SCPs to restrict resource creation that do not have the cost center and project ID on the resource.
D. Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources Update existing federated roles to restrict privileges to provision resources that do not include the cost center and project ID on the resource

Answer: B

NEW QUESTION 155
- (Exam Topic 2)
A solutions architect needs to provide AWS Cost and Usage Report data from a company's AWS Organizations management account The company already has an Amazon S3 bucket to store the reports The reports must be automatically ingested into a database that can be visualized with other toots.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE)

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that a new object creation in the S3 bucket will trigger
B. Create an AWS Cost and Usage Report configuration to deliver the data into the S3 bucket
C. Configure an AWS Glue crawler that a new object creation in the S3 bucket will trigger.
D. Create an AWS Lambda function that a new object creation in the S3 bucket will trigger
E. Create an AWS Glue crawler that me AWS Lambda function will trigger to crawl objects in me S3 bucket
F. Create an AWS Glue crawler that the Amazon EventBridge (Amazon CloudWatCh Events) rule will trigger to crawl objects m the S3 bucket

Answer: BDF

NEW QUESTION 157
- (Exam Topic 2)
A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances. Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances.
To meet regulatory and business requirements, the company must make the following changes for data backups:
• Backups must be retained based on custom daily, weekly, and monthly requirements.
• Backups must be replicated to at least one other AWS Region immediately after capture.
• The backup solution must provide a single source of backup status across the AWS environment.
• The backup solution must send immediate notifications upon failure of any resource backup.
Which combination of steps will meet these requirements with the LEAST amount of operational overhead? (Select THREE.)

A. Create an AWS Backup plan with a backup rule for each of the retention requirements.
B. Configure an AWS Backup plan to copy backups to another Region.
C. Create an AWS Lambda function to replicate backups to another Region and send notification if a failure occurs.
D. Add an Amazon Simple Notification Service (Amazon SNS) topic to the backup plan to send a notification for finished jobs that have any status except BACKUP_JOB_COMPLETEO.
E. Create an Amazon Data Lifecycle Manager (Amazon DLM) snapshot lifecycle policy for each of the retention requirements.
F. Set up RDS snapshots on each database.

Answer: BDE

NEW QUESTION 161
- (Exam Topic 2)
A company operates quick-service restaurants. The restaurants follow a predictable model with high sales traffic for -4 hours daily Sates traffic is lower outside of those peak hours.
The point of sale and management platform is deployed in the AWS Cloud and has a backend that is based or Amazon DynamoDB The database table uses provisioned throughput mode with 100.000 RCUs and 80.000 WCUs to match Known peak resource consumption.
The company wants to reduce its DynamoDB cost and minimize the operational overhead for the IT staff. Which solution meets these requirements MOST cost-effectively?

A. Reduce the provisioned RCUs and WCUs
B. Change the DynamoDB table to use on-demand capacity
C. Enable Dynamo DB auto seating for the table.
D. Purchase 1-year reserved capacity that is sufficient to cover the peak load for 4 hours each day.

Answer: C

NEW QUESTION 162
- (Exam Topic 2)
A company is planning to migrate an Amazon RDS for Oracle database to an RDS for PostgreSQL DB instance in another AWS account A solutions architect needs to design a migration strategy that will require no downtime and that will minimize the amount of time necessary to complete the migration The migration strategy must replicate all existing data and any new data that is created during the migration The target database must be identical to the source database at completion of the migration process
All applications currently use an Amazon Route 53 CNAME record as their endpoint for communication with the RDS for Oracle DB instance The RDS for Oracle DB instance is in a private subnet
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE)

A. Create a new RDS for PostgreSQL DB instance in the target account Use the AWS Schema Conversion Tool (AWS SCT) to migrate the database schema from the source database to the target database.
B. Use the AWS Schema Conversion Tool (AWS SCT) to create a new RDS for PostgreSQL DB instance in the target account with the schema and initial data from the source database
C. Configure VPC peering between the VPCs in the two AWS accounts to provide connectivity to both DB instances from the target accoun
D. Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account
E. Temporarily allow the source DB instance to be publicly accessible to provide connectivity from the VPC in the target account Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account.
F. Use AWS Database Migration Service (AWS DMS) in the target account to perform a full load plus change data capture (CDC) migration from the source database to the target database When the migration is complete, change the CNAME record to point to the target DB instance endpoint
G. Use AWS Database Migration Service (AWS DMS) in the target account to perform a change data capture (CDC) migration from the source database to the target database When the migration is complete change the CNAME record to point to the target DB instance endpoint

Answer: BCE

NEW QUESTION 165
- (Exam Topic 2)
A company owns a chain of travel agencies and is running an application in the AWS Cloud. Company employees use the application to search (or Information about travel destinations. Destination content is updated four times each year.
Two fixed Amazon EC2 instances serve the application. The company uses an Amazon Route 53 public hosted zone with a multivalue record of travel.example.com that returns the Elastic IP addresses for the EC2 instances. The application uses Amazon DynamoDB as its primary data store. The company uses a self-hosted Redis instance as a caching solution.
During content updates, the load on the EC2 instances and the caching solution increases drastically. This increased load has led to downtime on several

occasions. A solutions architect must update the application so that the application is highly available and can handle the load that is generated by the content updates.
Which solution will meet these requirements?

A. Set up DynamoDB Accelerator (DAX} as in-memory cach
B. Update the application to use DA
C. Create an Auto Scaling group for the EC2 instance
D. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the AL
E. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alia
F. Configure scheduled scaling for the EC2 instances before the content updates.
G. Set up Amazon ElastiCache for Redi
H. Update the application to use ElastiCach
I. Create an Auto Scaling group for the EC2 instance
J. Create an AmazonCloudFront distnbutio
K. and set the Auto Scaling group as an origin for the distributio
L. Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alias Manually scale up EC2 instances before the content updates
M. Set up Amazon ElastiCache for Memcache
N. Update the application to use ElastiCach
O. Create an Auto Scaling group for the EC2 instances Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the AL
P. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alia
Q. Configure scheduled scaling for the application before the content updates.
R. Set up DynamoDB Accelerator (DAX) as in-memory cach
S. Update the application to use DA
T. Create an Auto Scaling group for the EC2 instance
. Create an Amazon CloudFront distribution, and set the Auto Scaling group as an origin for the distributio
. Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alia
. Manually scale up EC2 instances before the content updates.

Answer: B

NEW QUESTION 167
- (Exam Topic 2)
A solutions architect is migrating an existing workload to AWS Fargate. The task can only run in a private subnet within the VPC where there is no direct connectivity from outside the system to the application When the Fargate task is launched the task fails with the following error:

```
CannotPullContainerError: API error (500): Get https://1111222233333.dkr.ecr.us-east-1.amazonaws.com/v2/: net/http: request canceled
while waiting for connection
```

How should the solutions architect correct this error?

A. Ensure the task is set to ENABLED for the auto-assign public IP setting when launching the task
B. Ensure the task is set to DISABLED (or the auto-assign public IP setting when launching the task Configure a NAT gateway in the public subnet in the VPC to route requests to the internet
C. Ensure the task is set to DISABLED for the auto-assign public IP setting when launching the task Configure a NAT gateway in the private subnet in the VPC to route requests to the internet
D. Ensure the network mode is set to bridge in the Fargate task definition

Answer: B

NEW QUESTION 172
- (Exam Topic 2)
A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode.
A recent RDS database failover test caused a 40-second outage to the application A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Use Amazon ElastiCache for Memcached in front of the database
B. Use Amazon ElastiCache for Redis in front of the database.
C. Use RDS Proxy in front of the database
D. Migrate the database to Amazon Aurora MySQL
E. Create an Amazon Aurora Replica
F. Create an RDS for MySQL read replica

Answer: ABF

NEW QUESTION 177
- (Exam Topic 2)
A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch.
Currently, the game consists of the following components deployed in a single AWS Region:
• Amazon S3 bucket that stores game assets
• Amazon DynamoDB table that stores player scores
A solutions architect needs to design a Region solution that wifi reduce latency improve reliability, and require the least effort to implement
What should the solutions architect do to meet these requirements'

A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket Configure S3Cross-Region Replication Create a new DynamoDB able in a new
Region Use the new table as a replica target tor DynamoDB global tables.
B. Create an Amazon CloudFront distribution to serve assets from the S3 bucke
C. Configure S3Same-Region Replicatio
D. Create a new DynamoDB able m a new Regio
E. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC)

F. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Regio
G. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
H. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets- Create an Amazon CloudFront distribution and configure origin failover with two origin accessing the S3 buckets Create a new DynamoDB table m a new Region Use the new table as a replica target for DynamoDB global tables.

Answer: B

NEW QUESTION 179
- (Exam Topic 2)
A car rental company has built a serverless REST API to provide data to its mobile app. The app consists of an Amazon API Gateway API with a Regional endpoint, AWS Lambda functions and an Amazon Aurora MySQL Serverless DB cluster The company recently opened the API to mobile apps of partners A significant increase in the number of requests resulted causing sporadic database memory errors Analysis of the API traffic indicates that clients are making multiple HTTP GET requests for the same queries in a short period of time Traffic is concentrated during business hours, with spikes around holidays and other events
The company needs to improve its ability to support the additional usage while minimizing the increase in costs associated with the solution.
Which strategy meets these requirements?

A. Convert the API Gateway Regional endpoint to an edge-optimized endpoint Enable caching in the production stage.
B. Implement an Amazon ElastiCache for Redis cache to store the results of the database calls Modify the Lambda functions to use the cache
C. Modify the Aurora Serverless DB cluster configuration to increase the maximum amount of available memory
D. Enable throttling in the API Gateway production stage Set the rate and burst values to limit the incoming calls

Answer: A

NEW QUESTION 183
- (Exam Topic 2)
A company runs applications on Amazon EC2 instances. The company plans to begin using an Auto Scaling group for the instances. As part of this transition, a solutions architect must ensure that Amazon CloudWatch Logs automatically collects logs from all new instances The new Auto Scaling group will use a launch template that includes the Amazon Linux 2 AMI and no key pair
Which solution meets these requirements?

A. Create an Amazon CloudWatch agent configuration for the workload Store the CloudWatch agent configuration in an Amazon S3 bucket Write an EC2 user data script to fetch the configuration He from Amazon S3. Configure the cloudWatch agent on the instance during Initial boot.
B. Create an Amazon CloudWatch agent configuration for the workload In AWS Systems Manager Parameter Store Create a Systems Manager document that Installs and configures the CloudWatch agent by using the configuration Create an Amazon EventBridge (Amazon CloudWatch Events) rule on the default event bus with a Systems Manager Run Command target that runs the document whenever an instance enters the running state.
C. Create an Amazon CloudWatch agent configuration for the workload Create an AWS Lambda function to Install and configure CloudWatch agent by using AWS Systems Manager Session Manage
D. Include the agent configuration inside the Lambda package Create an AWS Config custom rule to identify changes to the EC2 instances and invoke the Lambda function
E. Create an Amazon CloudWatch agent configuration for the workloa
F. Save the CloudWatch agent configuration as pan of an AWS Lambda deployment packag
G. Use AWS CloudTrail to capture EC2 tagging events and initiate agent installatio
H. Use AWS CodeBuild to configure the CloudWatch agent on the instances that run the workload.

Answer: B

NEW QUESTION 188
- (Exam Topic 2)
A company plans to refactor a monolithic application into a modern application designed deployed or AWS. The CLCD pipeline needs to be upgraded to support the modem design for the application with the following requirements
• It should allow changes to be released several times every hour.
* It should be able to roll back the changes as quickly as possible Which design will meet these requirements?

A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurations Deploy the application by replacing Amazon EC2 instances
B. Specify AWS Elastic Beanstak to sage in a secondary environment as the deployment target for the CI/CD pipeline of the applicatio
C. To deploy swap the staging and production environment URLs.
D. Use AWS Systems Manager to re-provision the infrastructure for each deployment Update the AmazonEC2 user data to pull the latest code art-fact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment
E. Roll out At application updates as pan of an Auto Scaling event using prebuilt AMI
F. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

Explanation:
It is the fastest when it comes to rollback and deploying changes every hour

NEW QUESTION 192
- (Exam Topic 2)
A company's site reliability engineer is performing a review of Amazon FSx for Windows File Server deployments within an account that the company acquired Company policy states that all Amazon FSx file systems must be configured to be highly available across Availability Zones.
During the review, the site reliability engineer discovers that one of the Amazon FSx file systems used a deployment type of Single-AZ 2 A solutions architect needs to minimize downtime while aligning this Amazon FSx file system with company policy.
What should the solutions architect do to meet these requirements?

A. Reconfigure the deployment type to Multi-AZ for this Amazon FSx tile system

B. Create a new Amazon FSx fie system with a deployment type o(Multi-A
C. Use AWS DataSync to transfer data to the new Amazon FSx file syste
D. Point users to the new location
E. Create a second Amazon FSx file system with a deployment type of Single-AZ 2. Use AWS DataSync to keep the data n syn
F. Switch users to the second Amazon FSx fie system in the event of failure
G. Use the AWS Management Console to take a backup of the Amazon FSx He system Create a new Amazon FSx file system with a deployment type of Multi-AZ Restore the backupto the new Amazon FSx file syste
H. Point users to the new location.

Answer: B

NEW QUESTION 195
- (Exam Topic 2)
A company is migrating an application to the AWS Cloud. The application runs in an on-premises data center and writes thousands of images into a mounted NFS file system each night After the company migrates the application, the company will host the application on an Amazon EC2 instance with a mounted Amazon Elastic File System (Amazon EFS) file system.
The company has established an AWS Direct Connect connection to AWS Before the migration cutover. a solutions architect must build a process that will replicate the newly created on-premises images to the EFS file system
What is the MOST operationally efficient way to replicate the images?

A. Configure a periodic process to run the aws s3 sync command from the on-premises file system to Amazon S3 Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system
B. Deploy an AWS Storage Gateway file gateway with an NFS mount poin
C. Mount the file gateway file system on the on-premises serve
D. Configure a process to periodically copy the images to the mount point
E. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system Send data over the Direct Connect connection to an S3 bucket by using a public VIF Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system
F. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system Send data over the Direct Connect connection to an AWS PrivateLink interface VPC endpoint for Amazon EFS by using a private VIF Configure a DataSync scheduled task to send the images to the EFS file system every 24 hours.

Answer: A

NEW QUESTION 196
- (Exam Topic 2)
A new startup is running a serverless application using AWS Lambda as the primary source of compute New versions of the application must be made available to a subset of users before deploying changes to all users Developers should also have the ability to stop the deployment and have access to an easy rollback mechanism A solutions architect decides to use AWS CodeDeploy to deploy changes when a new version is available.
Which CodeDeploy configuration should the solutions architect use?

A. A blue/green deployment
B. A linear deployment
C. A canary deployment
D. An all-at-once deployment

Answer: C

NEW QUESTION 201
- (Exam Topic 2)
A large education company recently introduced Amazon Workspaces to provide access to internal applications across multiple universities. The company is storing user proxies on an Amazon FSx for Windows File Server tile system. The Me system is configured with a DNS alias and is connected to a self-managed Active Directory As more users begin to use the Workspaces login time increases to unacceptable levels
An investigation reveals a degradation in performance of the file system. The company created the file system on HDD storage with a throughput of 16 MBps A solutions architect must improve the performance of the file system during a defined maintenance window
What should the solutions architect do to meet these requirements with the LEAST administrative effort?

A. Use AWS Backup to create a point-in-time backup of the file system Restore the backup to a new FSx for Windows File Server file system Select SSD as the storage type Select 32 MBps as the throughput capacity When the backup and restore process is completed adjust the DNS alias accordingly Delete the original file system
B. Disconnect users from the file system In the Amazon FSx console, update the throughput capacity to 32 MBps Update the storage type to SSD Reconnect users to the file system
C. Deploy an AWS DataSync agent onto a new Amazon EC2 instanc
D. Create a task Configure the existing file system as the source location Configure a new FSx forWindows File Server file system with SSD storage and 32 MBps of throughput as the target location Schedule the task When the task is completed adjust the DNS alias accordingly Delete the original file system.
E. Enable shadow copies on the existing file system by using a Windows PowerShell command Schedule the shadow copy job to create a point-in-time backup of the file system Choose to restore previousversions Create a new FSx for Windows File Server file system with SSD storage and 32 MBps of throughput When the copy job is completed, adjust the DNS alias Delete the original file system

Answer: D

NEW QUESTION 202
- (Exam Topic 2)
A company hosts its primary API on AWS by using an Amazon API Gateway API and AWS Lambda functions that contain the logic for the API methods. The company s internal applications use the API for core functionality and business logic. The company's customers use the API to access data from their accounts Several customers also have access to a legacy API that is running on a single standalone Amazon EC2 instance.
The company wants to increase the security for these APIs to better prevent denial of service (DoS) attacks, check for vulnerabilities, and guard against common exploits
What should a solutions architect do to meet these requirements?

A. Use AWS WAF to protect both APIs Configure Amazon Inspector to analyze the legacy API Configure Amazon GuardDuty to monitor for malicious attempts to access the APIs
B. Use AWS WAF to protect the API Gateway API Configure Amazon Inspector to analyze both APIs Configure Amazon GuardDuty to block malicious attempts to access the APIs.
C. Use AWS WAF to protect the API Gateway API Configure Amazon inspector to analyze the legacy APIConfigure Amazon GuardDuty to monitor for malicious attempts to access the APIs.
D. Use AWS WAF to protect the API Gateway API Configure Amazon inspector to protect the legacy API Configure Amazon GuardDuty to block malicious attempts to access the APIs.

Answer: C

NEW QUESTION 203
- (Exam Topic 2)
A solutions architect must update an application environment within AWS Elastic Beanstalk using a With green deployment methodology. The solutions architect creates an environment that is identical to the existing application environment and deploys the application to the new environment.
What should be done next to complete the update?

A. Redirect to the new environment using Amazon Route 53
B. Select the Swap Environment URLs option.
C. Replace the Auto Scaling launch configuration
D. Update the DNS records to point to the green environment

Answer: B

NEW QUESTION 208
- (Exam Topic 2)
A company has a platform that contains an Amazon S3 bucket for user content. The S3 bucket has thousands of terabytes of objects, all in the S3 Standard storage class. The company has an RTO of 6 hours The company must replicate the data from its primary AWS Region to a replication S3 bucket in another Region
The user content S3 bucket contains user-uploaded files such as videos and photos. The user content S3 bucket has an unpredictable access pattern. The number of users is increasing quickly, and the company wants to create an S3 Lifecycle policy to reduce storage costs
Which combination of steps will meet these requirements MOST cost-effectively'? (Select TWO)

A. Move the objects in the user content S3 bucket to S3 Intelligent-Tiering immediately
B. Move the objects in the user content S3 bucket to S3 Intelligent-Tiering after 30 days
C. Move the objects in the replication S3 bucket to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days and to S3 Glacier after 90 days
D. Move the objects in the replication S3 bucket to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days and to S3 Glacier Deep Archive after 90 days
E. Move the objects in the replication S3 bucket to S3 Standard-infrequent Access (S3 Standard-IA) after 30 days and to S3 Glacier Deep Archive after 180 days

Answer: AD

NEW QUESTION 213
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAP-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAP-C02 Product From:

https://www.2passeasy.com/dumps/SAP-C02/

Money Back Guarantee

SAP-C02 Practice Exam Features:

* SAP-C02 Questions and Answers Updated Frequently

* SAP-C02 Practice Questions Verified by Expert Senior Certified Staff

* SAP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SAP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year