Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
- (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 2
- (Exam Topic 1)
A company manages an on-premises JavaScript front-end web application. The application is hosted on two servers secured with a corporate Active Directory. The application calls a set of Java-based microservices on an application server and stores data in a clustered MySQL database. The application is heavily used during the day on weekdays. It is lightly used during the evenings and weekends.
Daytime traffic to the application has increased rapidly, and reliability has diminished as a result. The company wants to migrate the application to AWS with a solution that eliminates the need for server maintenance, with an API to securely connect to the microservices.
Which combination of actions will meet these requirements? (Select THREE.)

A. Host the web application on Amazon S3. Use Amazon Cognito identity pools (federated identities) with SAML for authentication and authorization.
B. Host the web application on Amazon EC2 with Auto Scalin
C. Use Amazon Cognito federation and Login with Amazon for authentication and authorization.
D. Create an API layer with Amazon API Gatewa
E. Rehost the microservices on AWS Fargate containers.
F. Create an API layer with Amazon API Gatewa
G. Rehost the microservices on Amazon Elastic Container Service (Amazon ECS) containers.
H. Replatform the database to Amazon RDS for MySQL.
I. Replatform the database to Amazon Aurora MySQL Serverless.

Answer: ACE

NEW QUESTION 3
- (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 4
- (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 5
- (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 am 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 Mocking 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 the correct IAM permissions to use the interface VPC endpoint.
F. Enable VPC Flow Logs and Amazon Detective Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic
G. Add an interface VPC endpoint for Kinesis Data Streams to the VP
H. 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 6
- (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 markeling 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 7
- (Exam Topic 1)
A company needs to create and manage multiple AWS accounts for a number of departments from a central location. The security team requires read-only access to all accounts from its own AWs account. The company is using AWS Organizations and created an account tor the security team.
How should a solutions architect meet these requirements?

A. Use the OrganizationAccountAccessRole IAM role to create a new IAM policy wilh read-only access in each member accoun
B. Establish a trust relationship between the IAM policy in each member account and the security accoun
C. Ask the security team lo use the IAM policy to gain access.
D. Use the OrganizationAccountAccessRole IAM role to create a new IAM role with read-only access in each member accoun
E. Establish a trust relationship between the IAM role in each member account and the security accoun
F. Ask the security team lo use the IAM role to gain access.
G. Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the OrganizationAccountAccessRole IAM role in the master account from the security accoun
H. Use the generated temporary credentials to gain access.
I. Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the OrganizationAccountAccessRole IAM role in the member account from the security accoun

J. Use the generated temporary credentials to gain access.

Answer: D

NEW QUESTION 8
- (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 9
- (Exam Topic 1)
A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency network connections between all of the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant.
Which solution will meet these requirements?

A. Launch five new EC2 instances into a cluster placement grou
B. Ensure that the EC2 instance type supports enhanced networking.
C. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zon
D. Attach an extra elastic network interface to each EC2 instance.
E. Launch five new EC2 instances into a partition placement grou
F. Ensure that the EC2 instance type supports enhanced networking.
G. Launch five new EC2 instances into a spread placement grou
H. Attach an extra elastic network interface to each EC2 instance.

Answer: A

Explanation:
When you launch EC2 instances in a cluster they benefit from performance and low latency. No redundancy though as per the question
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html.

NEW QUESTION 10
- (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 10
- (Exam Topic 1)
A company 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 13
- (Exam Topic 1)
A solutions architect is designing a network for a new cloud deployment. Each account will need autonomy to modify route tables and make changes. Centralized and controlled egress internet connectivity is also needed. The cloud footprint is expected to grow to thousands of AWS accounts.
Which architecture will meet these requirements?

A. A centralized transit VPC with a VPN connection to a standalone VPC in each accoun
B. Outbound internet traffic will be controlled by firewall appliances.
C. A centralized shared VPC with a subnet for each accoun
D. Outbound internet traffic will controlled through a fleet of proxy servers.
E. A shared services VPC to host central assets to include a fleet of firewalls with a route to the internet.Each spoke VPC will peer to the central VPC.
F. A shared transit gateway to which each VPC will be attache
G. Outbound internet access will route through a fleet of VPN-attached firewalls.

Answer: D

Explanation:
https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/centr
https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/centr
AWS Transit Gateway helps you design and implement networks at scale by acting as a cloud router. As your network grows, the complexity of managing incremental connections can slow you down. AWS Transit Gateway connects VPCs and on-premises networks through a central hub. This simplifies your network and puts an end to complex peering relationships -- each new connection is only made once.

NEW QUESTION 15
- (Exam Topic 1)
A company has 50 AWS accounts that are members of an organization in AWS Organizations Each account contains multiple VPCs The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.
Which combination of steps will meet these requirements? (Select TWO)

A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager
B. Prom the management account, share the transit gateway with member accounts by using an AWSOrganizations SCP
C. Launch an AWS CloudFormation stack set from the management account that automatical^/ creates a new VPC and a VPC transit gateway attachment in a member accoun
D. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.
E. Launch an AWS CloudFormation stack set from the management account that automatical^ creates a new VPC and a peering transit gateway attachment in a member accoun
F. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.
G. From the management account, share the transit gateway with member accounts by using AWS Service Catalog

Answer: AC

Explanation:
https://aws.amazon.com/blogs/mt/self-service-vpcs-in-aws-control-tower-using-aws-service-catalog/

NEW QUESTION 16
- (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 17
- (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 18
- (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 20
- (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 23
- (Exam Topic 1)
A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand.
Which solutions meet these requirements? (Choose two.)

A. Create an Amazon API Gateway REST AP
B. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.
C. Create an Amazon API Gateway HTTP AP
D. Configure this API with direct integrations to Dynamo DB by using API Gateway's AWS integration type.
E. Create an Amazon API Gateway HTTP AP
F. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.
G. Create an accelerator in AWS Global Accelerato
H. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.
I. Create a Network Load Balance
J. Configure listener rules to forward requests to the appropriate AWS Lambda functions

Answer: CD

Explanation:

https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-dynamo-db.html

NEW QUESTION 27
- (Exam Topic 1)
A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services. Database credentials are hard-coded on each instance. SSH keys for command-tine remote access are stored in a secured Amazon S3 bucket. The company has asked its solutions architect to improve the security posture of the architecture without adding operational complexity.
Which combination of steps should the solutions architect take to accomplish this? (Select THREE.)

A. Use Amazon EC2 instance profiles with an IAM role.
B. Use AWS Secrets Manager to store access keys and secret access keys.
C. Use AWS Systems Manager Parameter Store to store database credentials.
D. Use a secure fleet of Amazon EC2 bastion hosts (or remote access.
E. Use AWS KMS to store database credentials.
F. Use AWS Systems Manager Session Manager tor remote access

Answer: ACF

Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html

NEW QUESTION 32
- (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 33
- (Exam Topic 1)
A company uses an on-premises data analytics platform. The system is highly available in a fully redundant configuration across 12 servers in the company's data center.
The system runs scheduled jobs, both hourly and daily, in addition to one-time requests from users. Scheduled jobs can take between 20 minutes and 2 hours to finish running and have tight SLAs. The scheduled jobs account for 65% of the system usage. User jobs typically finish running in less than 5 minutes and have no SLA. The user jobs account for 35% of system usage. During system failures, scheduled jobs must continue to meet SLAs. However, user jobs can be delayed.
A solutions architect needs to move the system to Amazon EC2 instances and adopt a consumption-based model to reduce costs with no long-term commitments. The solution must maintain high availability and must not affect the SLAs.
Which solution will meet these requirements MOST cost-effectively?

A. Split the 12 instances across two Availability Zones in the chosen AWS Regio
B. Run two instances in each Availability Zone as On-Demand Instances with Capacity Reservation
C. Run four instances in each Availability Zone as Spot Instances.
D. Split the 12 instances across three Availability Zones in the chosen AWS Regio
E. In one of the Availability Zones, run all four instances as On-Demand Instances with Capacity Reservation
F. Run the remaining instances as Spot Instances.
G. Split the 12 instances across three Availability Zones in the chosen AWS Regio
H. Run two instances in each Availability Zone as On-Demand Instances with a Savings Pla
I. Run two instances in each Availability Zone as Spot Instances.
J. Split the 12 instances across three Availability Zones in the chosen AWS Regio
K. Run three instances in each Availability Zone as On-Demand Instances with Capacity Reservation
L. Run one instance in each Availability Zone as a Spot Instance.

Answer: D

NEW QUESTION 35
- (Exam Topic 1)
A solutions architect is designing an application to accept timesheet entries from employees on their mobile devices. Timesheets will be submitted weekly, with most of the submissions occurring on Friday. The data must be stored in a format that allows payroll administrators to run monthly reports. The infrastructure must be highly available and scale to match the rate of incoming data and reporting requests.
Which combination of steps meets these requirements while minimizing operational overhead? (Select TWO.)

A. Deploy the application to Amazon EC2 On-Demand Instances With load balancing across multiple Availability Zone
B. Use scheduled Amazon EC2 Auto Scaling to add capacity before the high volume of submissions on Fridays.
C. Deploy the application in a container using Amazon Elastic Container Service (Amazon ECS) with load balancing across multiple Availability Zone
D. Use scheduled Service Auto Scaling to add capacity before the high volume of submissions on Fridays.
E. Deploy the application front end to an Amazon S3 bucket served by Amazon CloudFron

F. Deploy the application backend using Amazon API Gateway with an AWS Lambda proxy integration.
G. Store the timesheet submission data in Amazon Redshif
H. Use Amazon QuickSight to generate the reports using Amazon Redshift as the data source.
I. Store the timesheet submission data in Amazon S3. Use Amazon Athena and Amazon QuickSight to generate the reports using Amazon S3 as the data source.

Answer: AE

NEW QUESTION 40
- (Exam Topic 1)
A company that tracks medical devices in hospitals wants to migrate its existing storage solution to the AWS Cloud. The company equips all of its devices with sensors that collect location and usage information. This sensor data is sent in unpredictable patterns with large spikes. The data is stored in a MySQL database running on premises at each hospital. The company wants the cloud storage solution to scale with usage.
The company's analytics team uses the sensor data to calculate usage by device type and hospital. The team needs to keep analysis tools running locally while fetching data from the cloud. The team also needs to use existing Java application and SQL queries with as few changes as possible.
How should a solutions architect meet these requirements while ensuring the sensor data is secure?

A. Store the data in an Amazon Aurora Serverless databas
B. Serve the data through a Network Load Balancer (NLB). Authenticate users using the NLB with credentials stored in AWS Secrets Manager.
C. Store the data in an Amazon S3 bucke
D. Serve the data through Amazon QuickSight using an IAM user authorized with AWS Identity and Access Management (IAM) with the S3 bucket as the data source.
E. Store the data in an Amazon Aurora Serverless databas
F. Serve the data through the Aurora Data API using an IAM user authorized with AWS Identity and Access Management (IAM) and the AWS Secrets Manager ARN.
G. Store the data in an Amazon S3 bucke
H. Serve the data through Amazon Athena using AWS PrivateLink to secure the data in transit.

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/new-data-api-for-amazon-aurora-serverless/ https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html
https://aws.amazon.com/blogs/aws/aws-privatelink-for-amazon-s3-now-available/ https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html#data-api.access
The data is currently stored in a MySQL database running on-prem. Storing MySQL data in S3 doesn't sound good so B & D are out. Aurora Data API "enables the SQL HTTP endpoint, a connectionless Web Service API for running SQL queries against this database. When the SQL HTTP endpoint is enabled, you can also query your database from inside the RDS console (these features are free to use)."

NEW QUESTION 41
- (Exam Topic 1)
A company has many AWS accounts and uses AWS Organizations to manage all of them. A solutions architect must implement a solution that the company can use to share a common network across multiple accounts.
The company's infrastructure team has a dedicated infrastructure account that has a VPC. The infrastructure team must use this account to manage the network. Individual accounts cannot have the ability to manage their own networks. However, individual accounts must be able to create AWS resources within subnets.
Which combination of actions should the solutions architect perform to meet these requirements? (Select TWO.)

A. Create a transit gateway in the infrastructure account.
B. Enable resource sharing from the AWS Organizations management account.
C. Create VPCs in each AWS account within the organization in AWS Organization
D. Configure the VPCs to share the same CIDR range and subnets as the VPC in the infrastructure accoun
E. Peer the VPCs in each individual account with the VPC in the infrastructure account,
F. Create a resource share in AWS Resource Access Manager in the infrastructure accoun
G. Select the specific AWS Organizations OU that will use the shared networ
H. Select each subnet to associate with the resource share.
I. Create a resource share in AWS Resource Access Manager in the infrastructure accoun
J. Select the specific AWS Organizations OU that will use the shared networ
K. Select each prefix list to associate with the resource share.

Answer: CE

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html

NEW QUESTION 46
- (Exam Topic 1)
A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics After extensive testing, the company determines that the m5.2xlarge instance size is optimal for the workload Application data is stored in db.r4.4xlarge Amazon RDS instances that are confirmed to be optimal. The traffic to the web application spikes randomly during the day.
What other cost-optimization methods should the company implement to further reduce costs without impacting the reliability of the application?

A. Double the instance count in the Auto Scaling groups and reduce the instance size to m5.large
B. Reserve capacity for the RDS database and the minimum number of EC2 instances that are constantly running.
C. Reduce the RDS instance size to db.r4.xlarge and add five equivalent^ sized read replicas to provide reliability.
D. Reserve capacity for all EC2 instances and leverage Spot Instance pricing for the RDS database.

Answer: B

Explanation:
People are being confused by the term 'reserve capacity'. This is not the same as an on-demand capacity reservation. This article by AWS clearly states that by 'reserving capacity' you are reserving the instances and reducing your costs. See https://aws.amazon.com/aws-cost-management/aws-cost-optimization/reserved-instances/

NEW QUESTION 48
- (Exam Topic 1)
A company hosts a photography website on AWS that has global visitors. The website has experienced steady increases in traffic during the last 12 months, and users have reported a delay in displaying images. The company wants to configure Amazon CloudFront lo deliver photos to visitors with minimal latency.
Which actions will achieve this goal? (Select TWO.)

A. Set the Minimum TTL and Maximum TTL to 0 in the CloudFront distribution.
B. Set the Minimum TTL and Maximum TTL to a high value in the CloudFront distribution.
C. Set the CloudFront distribution to forward all headers, all cookies, and all query strings to the origin.
D. Set up additional origin servers that are geographically closer to the requester
E. Configure latency-based routing in Amazon Route 53.
F. Select Price Class 100 on Ihe CloudFront distribution.

Answer: BD

NEW QUESTION 49
- (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 50
- (Exam Topic 1)
A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer The security team requires that all application access attempts be made available for analysis Information about the client IP address, connection type, and user agent must be included.
Which solution will meet these requirements?

A. Enable EC2 detailed monitoring, and include network logs Send all logs through Amazon Kinesis Data Firehose to an Amazon ElasDcsearch Service (Amazon ES) cluster that the security team uses for analysis.
B. Enable VPC Flow Logs for all EC2 instance network interfaces Publish VPC Flow Logs to an Amazon S3 bucket Have the security team use Amazon Athena to query and analyze the logs
C. Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket Have the security team use Amazon Athena to query and analyze the logs
D. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the sourc
E. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elastic search Service (Amazon ES) cluster that the security team uses for analysis.

Answer: C

Explanation:
https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html

NEW QUESTION 53
- (Exam Topic 1)
A company is using AWS CodePipeline for the CI/CO of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts.
As the application has become more complex, recent resource changes in the Cloud Formation templates have caused unplanned downtime.
How should a solutions architect improve the CI'CD pipeline to reduce the likelihood that changes in the templates will cause downtime?

A. Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployment
B. Write test plans for a testing team to execute in a non-production environment before approving the change for production.
C. Implement automated testing using AWS CodeBuild in a test environmen

D. Use CloudFormation change sets to evaluate changes before deploymen

E. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.

F. Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correc

G. Adapt the deployment code to check for error conditions and generate notifications on error

H. Deploy to a test environment and execute a manual test plan before approving the change for production.

I. Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment script

J. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.

Answer: B

Explanation:
https://aws.amazon.com/blogs/devops/performing-bluegreen-deployments-with-aws-codedeploy-and-auto-scalin When one adopts go infrastructure as code, we need to test the infrastructure code as well via automated testing, and revert to original if things are not performing correctly.

NEW QUESTION 55
- (Exam Topic 1)
A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes lo the items stored in the DynamoDB tables must be logged within 30 minutes.
Which solution meets the requirements?

A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them (or anomalous behavior

B. Send Amazon SNS notifications when anomalous behaviors are detected.

C. Use AWS CloudTrail to capture all the APIs that change the DynamoDB table

D. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.

E. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambd

F. Create a Lambda function to output records lo Amazon Kinesis Data Stream

G. Analyze any anomalies with Amazon Kinesis Data Analytic

H. Send SNS notifications when anomalous behaviors are detected.

I. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda (unction as a target to analyze behavio

J. Send SNS notifications when anomalous behaviors are detected.

Answer: C

Explanation:
https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/#:~:text=DynamoDB DynamoDb Stream to capture DynamoDB update. And Kinesis Data Analytics for anomaly detection (it uses
AWS proprietary Random Cut Forest Algorithm)

NEW QUESTION 57
- (Exam Topic 1)
A company has a website that enables users to upload videos. Company policy states the uploaded videos must be analyzed for restricted content. An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SOS queue with the video's location. A backend application pulls this location from Amazon SOS and analyzes the video.
The video analysis is compute-intensive and occurs sporadically during the day The website scales with demand. The video analysis application runs on a fixed number of instances. Peak demand occurs during the holidays, so the company must add instances to the application dunng this time. All instances used are currently on-demand Amazon EC2 T2 instances. The company wants to reduce the cost of the current solution.
Which of the following solutions is MOST cost-effective?

A. Keep the website on T2 instance

B. Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to cover peak deman

C. Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application

D. Keep the website on T2 instance

E. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak deman

F. Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances.

G. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instance

H. Determine the minimum number of website instances required during off-peak times and use On-Demand Instances to cover them while using Spot capacity to cover peak demand Use Spot Fleet for the video anarysis application comprised of C4 and Amazon EC2 C5 instances.

I. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instance

J. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand Use Spot Fleet for the video analysis application comprised of R4 and Amazon EC2 R5 instances

Answer: B

NEW QUESTION 62
- (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 theVPC peering connection.

L. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized 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 64
- (Exam Topic 1)
A scientific organization requires the processing of text and picture data stored in an Amazon S3 bucket. The data is gathered from numerous radar stations during a mission's live, time-critical phase. The data is uploaded by the radar stations to the source S3 bucket. The data is preceded with the identification number of the radar station.
In a second account, the business built a destination S3 bucket. To satisfy a compliance target, data must be transferred from the source S3 bucket to the destination S3 bucket. Replication is accomplished by using an S3 replication rule that covers all items in the source S3 bucket.
A single radar station has been recognized as having the most precise data. At this radar station, data replication must be completed within 30 minutes of the radar station uploading the items to the source S3 bucket.
What actions should a solutions architect take to ensure that these criteria are met?

A. Set up an AWS DataSync agent to replicate the prefixed data from the source S3 bucket to the destination S3 bucke

B. Select to use at available bandwidth on the task, and monitor the task to ensure that it is in the TRANSFERRING statu

C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.

D. In the second account, create another S3 bucket to receive data from the radar station with the most accurate data Set up a new replication rule for this new S3 bucket to separate the replication from the other radar stations Monitor the maximum replication time to the destinatio

E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold

F. Enable Amazon S3 Transfer Acceleration on the source S3 bucket, and configure the radar station with the most accurate data to use the new endpoint Monitor the S3 destination bucket's TotalRequestLatency metric Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes

G. Create a new S3 replication rule on the source S3 bucket that filters for the keys that use the prefix of the radar station with the most accurate data Enable S3 Replication Time Control (S3 RTC) Monitor the maximum replication time to the destination Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-time-control.html

NEW QUESTION 67
- (Exam Topic 1)
A company is planning to set up a REST API application on AWS. The application team wants to set up a new identity store on AWS The IT team does not want to maintain any infrastructure or servers for this deployment.
What is the MOST operationally efficient solution that meets these requirements?

A. Deploy the application as AWS Lambda function

B. Set up Amazon API Gateway REST API endpoints for the application Create a Lambda function, and configure a Lambda authorizer

C. Deploy the application in AWS AppSync, and configure AWS Lambda resolvers Set up an Amazon Cognito user pool, and configure AWS AppSync to use the user pool for authorization

D. Deploy the application as AWS Lambda function

E. Set up Amazon API Gateway REST API endpoints for the application Set up an Amazon Cognito user pool, and configure an Amazon Cognito authorizer

F. Deploy the application in Amazon Elastic Kubemetes Service (Amazon EKS) cluster

G. Set up an Application Load Balancer for the EKS pods Set up an Amazon Cognito user pool and service pod for authentication.

Answer: C

NEW QUESTION 70
- (Exam Topic 1)
A company runs a popular public-facing ecommerce website. Its user base is growing quickly from a local market to a national market. The website is hosted in an on-premises data center with web servers and a MySQL database. The company wants to migrate its workload (o AWS. A solutions architect needs to create a solution to:
• Improve security
• Improve reliability Improve availability
• Reduce latency
• Reduce maintenance
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Use Amazon EC2 instances in two Availability Zones for the web servers in an Auto Scaling group behind an Application Load Balancer.

B. Migrate the database to a Multi-AZ Amazon Aurora MySQL DB cluster.

C. Use Amazon EC2 instances in two Availability Zones to host a highly available MySQL database cluster.

D. Host static website content in Amazon S3. Use S3 Transfer Acceleration to reduce latency while serving webpage

E. Use AWS WAF to improve website security.

F. Host static website content in Amazon S3. Use Amazon CloudFronl to reduce latency while serving webpage

G. Use AWS WAF to improve website security

H. Migrate the database to a single-AZ Amazon RDS for MySQL DB instance.

Answer: ABE

NEW QUESTION 72
- (Exam Topic 1)
A company runs an application on AWS. An AWS Lambda function uses credentials to authenticate to an Amazon RDS tor MySQL DB instance. A security risk assessment identified that these credentials are not frequently rotated. Also, encryption at rest is not enabled for the DB instance. The security team requires that both of these issues be resolved.
Which strategy should a solutions architect recommend to remediate these security risks?

A. Configure the Lambda function to store and retrieve the database credentials in AWS Secrets Manager and enable rotation of the credential
B. Take a snapshot ol the DB instance and encrypt a copy of that snapsho
C. Replace the DB instance with a new DB instance that is based on the encrypted snapshot.
D. Enable IAM DB authentication on the DB instanc
E. Grant the Lambda execution role access to the DB instanc
F. Modify the DB instance and enable encryption.
G. Enable IAM DB authentication on the DB instanc
H. Grant the Lambda execution role access to the DB instanc
I. Create an encrypted read replica of the DB instanc
J. Promote Ihe encrypted read replica to be the new primary node.
K. Configure the Lambda function to store and retrieve the database credentials as encrypted AWS Systems Manager Parameter Store parameter
L. Create another Lambda function to automatically rotate the credential
M. Create an encrypted read replica of the DB instanc
N. Promote the encrypted read replica to be the new primary node.

Answer: A

Explanation:
Parameter store can store DB credentials as secure string but CANNOT rotate secrets, hence, go with A + Cannot enable encryption on existing MySQL RDS instance, must create a new encrypted one from unencrypted snapshot.
https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets- Encrypting a unencrypted instance of DB or creating a encrypted replica of an un encrypted DB instance are not possible Hence A is the only solution possible.
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption.

NEW QUESTION 73
- (Exam Topic 1)
A company runs an e-commerce platform with front-end and e-commerce tiers. Both tiers run on LAMP stacks with the front-end instances running behind a load balancing appliance that has a virtual offering on AWS Current*/, the operations team uses SSH to log in to the instances to maintain patches and address other concerns. The platform has recently been the target of multiple attacks, including.
• A DDoS attack.
• An SOL injection attack
• Several successful dictionary attacks on SSH accounts on the web servers
The company wants to improve the security of the e-commerce platform by migrating to AWS. The company's solutions architects have decided to use the following approach;
• Code review the existing application and fix any SQL injection issues.
• Migrate the web application to AWS and leverage the latest AWS Linux AMI to address initial security patching.
• Install AWS Systems Manager to manage patching and allow the system administrators to run commands on all instances, as needed.
What additional steps will address all of the identified attack types while providing high availability and minimizing risk?

A. Enable SSH access to the Amazon EC2 instances using a security group that limits access to specific IP
B. Migrate on-premises MySQL to Amazon RDS Multi-AZ Install the third-party load balancer from the AWS Marketplace and migrate the existing rules to the load balancer's AWS instances Enable AWS Shield Standard for DDoS protection
C. Disable SSH access to the Amazon EC2 instance
D. Migrate on-premises MySQL to Amazon RDS Multi-AZ Leverage an Elastic Load Balancer to spread the load and enable AWS Shield Advanced for protectio
E. Add an Amazon CloudFront distribution in front of the website Enable AWS WAF on the distribution to manage the rules.
F. Enable SSH access to the Amazon EC2 instances through a bastion host secured by limiting access to specific IP addresse
G. Migrate on-premises MySQL to a self-managed EC2 instanc
H. Leverage an AWS Elastic Load Balancer to spread the load, and enable AWS Shield Standard for DDoS protection Add anAmazon CloudFront distribution in front of the website.
I. Disable SSH access to the EC2 instance
J. Migrate on-premises MySQL to Amazon RDS Single-A
K. Leverage an AWS Elastic Load Balancer to spread the load Add an Amazon CloudFront distribution in front of the website Enable AWS WAF on the distribution to manage the rules.

Answer: B

NEW QUESTION 76
- (Exam Topic 1)
A company has implemented an ordering system using an event-dnven architecture. Dunng initial testing, the system stopped processing orders Further tog analysis revealed that one order message in an Amazon Simple Queue Service (Amazon SOS) standard queue was causing an error on the backend and blocking all subsequent order messages The visibility timeout of the queue is set to 30 seconds, and the backend processing timeout is set to 10 seconds. A solutions architect needs to analyze faulty order messages and ensure that the system continues to process subsequent messages
Which step should the solutions architect take to meet these requirements?

A. Increase the backend processing timeout to 30 seconds to match the visibility timeout
B. Reduce the visibility timeout of the queue to automatically remove the faulty message
C. Configure a new SOS FIFO queue as a dead-letter queue to isolate the faulty messages
D. Configure a new SOS standard queue as a dead-letter queue to isolate the faulty messages.

Answer: D

NEW QUESTION 79
- (Exam Topic 1)
A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The maximum size ol the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:
VPC CIDR: 10.0.0.0/23
AZ1 subnet CIDR: 10.0.0.0/24 AZ2 subnet CIDR: 10.0.1.0/24
Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime.
Which solution will meet these requirements?

A. Update the Auto Scaling group to use the AZ2 subnet onl
B. Delete and re-create the AZ1 subnet using hall the previous address spac
C. Adjust the Auto Seating group to also use the new AZ1 subne
D. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet onl
E. Remove the current AZ2 subne
F. Create a new AZ2 subnet using the second half of the address space from the original AZ1 subne
G. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.
H. Terminate the EC2 instances in the AZ1 subne
I. Delete and re-create the AZ1 subnet using half the address spac
J. Update the Auto Scaling group to use this new subne
K. Repeat this for the second A
L. Define a new subnet in AZ3, then update the Auto Scaling group to target all three new subnets.
M. Create a new VPC with the same IPv4 address space and define three subnets, with one for each A
N. Update the existing Auto Scaling group to target the new subnets in the new VPC.
O. Update the Auto Scaling group to use the AZ2 subnet onl
P. Update the AZ1 subnet to have half the previous address spac
Q. Adjust the Auto Scaling group to also use the AZ1 subnet agai
R. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet onl
S. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subne
T. Create a new AZ3 subnet using halt the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/vpc-ip-address-range/?nc1=h_ls
It's not possible to modify the IP address range of an existing virtual private cloud (VPC) or subnet. You must delete the VPC or subnet, and then create a new VPC or subnet with your preferred CIDR block.

NEW QUESTION 84
- (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 85
- (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 88
- (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 91
- (Exam Topic 1)
A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.
Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

A. Enable Aurora Auto Scaling for Aurora Replica
B. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.
C. Enable Aurora Auto Scaling for Aurora writer
D. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
E. Enable Aurora Auto Scaling for Aurora Replica
F. Use an Application Load Balancer with the round robin routing and sticky sessions enabled.
G. Enable Aurora Scaling for Aurora writer
H. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

Explanation:
Aurora Auto Scaling enables your Aurora DB cluster to handle sudden increases in connectivity or workload. When the connectivity or workload decreases, Aurora Auto Scaling removes unnecessary Aurora Replicas so that you don't pay for unused provisioned DB instances

NEW QUESTION 93
- (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: A

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

NEW QUESTION 95
- (Exam Topic 1)
A company has a policy that all Amazon EC2 instances that are running a database must exist within the same subnets in a shared VPC Administrators must follow security compliance requirements and are not allowed to directly log in to the shared account All company accounts are members of the same organization in AWS Organizations. The number of accounts will rapidly increase as the company grows.
A solutions architect uses AWS Resource Access Manager to create a resource share in the shared account What is the MOST operationally efficient configuration to meet these requirements?

A. Add the VPC to the resource shar
B. Add the account IDs as principals
C. Add all subnets within the VPC to the resource shar
D. Add the account IDs as principals
E. Add all subnets within the VPC to the resource shar
F. Add the organization as a principal.
G. Add the VPC to the resource shar
H. Add the organization as a principal

Answer: C

Explanation:
https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html#getting-started-sharing-create To restrict resource sharing to only principals in your organization, choose Allow sharing with principals in your organization only.
https://docs.aws.amazon.com/ram/latest/userguide/ram-ug.pdf

NEW QUESTION 100
- (Exam Topic 1)
To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network. The security team mandates that no traffic accessing this data should traverse the public internet.
How should the solutions architect design a highly available solution that meets the requirements and is cost-effective?

A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use.Use the company WAN lo send traffic over to the headquarters and then to the respective DXconnection to access the data.
B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region.Use the company WAN to send traffic over a DX connectio
C. Use inter-region VPC peering to access the data in other AWS Regions.
D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region.Use the company WAN to send traffic over a DX connectio
E. Use an AWS transit VPC solution to access data in other AWS Regions.
F. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region.Use the company WAN to send traffic over a DX connectio
G. Use Direct Connect Gateway to access data in other AWS Regions.

Answer: D

Explanation:
This feature also allows you to connect to any of the participating VPCs from any Direct Connect location, further reducing your costs for making using AWS services on a cross-region basis. https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/
https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-g

NEW QUESTION 101
- (Exam Topic 1)
A company wants to migrate its corporate data center from on premises to the AWS Cloud. The data center includes physical servers and VMs that use VMware and Hyper-V. An administrator needs to select the correct services to collect data (or the initial migration discovery process. The data format should be supported by AWS Migration Hub. The company also needs the ability to generate reports from the data.
Which solution meets these requirements?

A. Use the AWS Agentless Discovery Connector for data collection on physical servers and all VM
B. Store the collected data in Amazon S3. Query the data with S3 Selec
C. Generate reports by using Kibana hosted on Amazon EC2.
D. Use the AWS Application Discovery Service agent for data collection on physical servers and all VMs.Store the collected data in Amazon Elastic File System (Amazon EFS). Query the data and generate reports with Amazon Athena.
E. Use the AWS Application Discovery Service agent for data collection on physical servers and Hyper-
F. Use the AWS Agentless Discovery Connector for data collection on VMwar
G. Store the collected data in Amazon S3. Query the data with Amazon Athen
H. Generate reports by using Amazon QuickSight.
I. Use the AWS Systems Manager agent for data collection on physical server
J. Use the AWS Agentless Discovery Connector for data collection on all VM
K. Store, query, and generate reports from the collected data by using Amazon Redshift.

Answer: C

Explanation:
https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-connector.html

NEW QUESTION 104

- (Exam Topic 1)
A company is storing data on premises on a Windows file server. The company produces 5 GB of new data
daily.
The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS.
Which data migration strategy should the company use?

A. Use the file gateway option in AWS Storage Gateway to replace the existing Windows file server, and point the existing file share to the new file gateway.
B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.
C. Use AWS Data Pipeline to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).
D. Use AWS DataSync to schedule a daily task lo replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS),

Answer: B

Explanation:
https://aws.amazon.com/storagegateway/file/ https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html
https://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os-win

NEW QUESTION 109
- (Exam Topic 1)
A multimedia company needs to deliver its video-on-demand (VOD) content to its subscribers in a
cost-effective way. The video files range in size from 1-15 GB and are typically viewed frequently for the first 6 months alter creation, and then access decreases considerably. The company requires all video files to remain immediately available for subscribers. There are now roughly 30.000 files, and the company anticipates doubling that number over time.
What is the MOST cost-effective solution for delivering the company's VOD content?

A. Store the video files in an Amazon S3 bucket using S3 Intelligent-Tierin
B. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.
C. Use AWS Elemental MediaConvert and store the adaptive bitrate video files in Amazon S3. Configure an AWS Elemental MediaPackage endpoint to deliver the content from Amazon S3.
D. Store the video files in Amazon Elastic File System (Amazon EFS) Standar
E. Enable EFS lifecycle management to move the video files to EFS Infrequent Access after 6 month
F. Create an Amazon EC2 Auto Scaling group behind an Elastic Load Balancer to deliver the content from Amazon EFS.
G. Store the video files in Amazon S3 Standar
H. Create S3 Lifecycle rules to move the video files to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months and to S3 Glacier Deep Archive after 1 yea
I. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.

Answer: A

Explanation:
https://d1.awsstatic.com/whitepapers/amazon-cloudfront-for-media.pdf https://aws.amazon.com/solutions/implementations/video-on-demand-on-aws/

NEW QUESTION 110
- (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 (ALBJ and put the read replicas behind the ALB.
E. Configure an Amazon CloudWatch alarm to detect a failed read replica Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.
F. 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 111
- (Exam Topic 1)
A solutions architect must analyze a company's Amazon EC2 Instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently The company is running several large, high-memory EC2 instances lo host database dusters that are deployed in active/passive configurations The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern The solutions architect must analyze the environment and take action based on the findings. Which solution meets these requirements MOST cost-effectively?

A. Create a dashboard by using AWS Systems Manager OpsConter Configure visualizations tor Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes Review the dashboard periodically and identify usage patterns Right size the EC2 instances based on the peaks in the metrics
B. Turn on Amazon CloudWatch detailed monitoring for the EC2 instances and their EBS volumes Create and review a dashboard that is based on the metrics Identify usage patterns Right size the FC? instances based on the peaks In the metrics
C. Install the Amazon CloudWatch agent on each of the EC2 Instances Turn on AWS Compute Optimizer, and let it run for at least 12 hours Review the recommendations from Compute Optimizer, and right size the EC2 instances as directed
D. Sign up for the AWS Enterprise Support plan Turn on AWS Trusted Advisor Wait 12 hours Review the recommendations from Trusted Advisor, and rightsize the EC2 instances as directed

Answer: C

Explanation:
(https://aws.amazon.com/compute-optimizer/pricing/ , https://aws.amazon.com/systems-manager/pricing/). https://aws.amazon.com/compute-optimizer/

NEW QUESTION 116
- (Exam Topic 2)
A company wants to improve cost awareness for its Amazon EMR platform The company has aWocated budgets for each team's Amazon EMR usage When a budgetary threshold is reached a notification should be sent by email to the budget office's distribution list Teams should be able lo view their EMR cluster expenses to date A solutions architect needs to create a solution that ensures this policy is proactively and centrally enforced in a multi-account environment Which combination of steps should the solutions architect take to meet these requirements? (Select TWO.)

A. Update the AWS CloudFormation template to include the AWS Budgets Budget resource with the NotificationsWithSubscnbers property
B. Implement Amazon CloudWatch dashboards for Amazon EMR usage
C. Create an EMR bootstrap action that runs at startup that calls the Cost Explorer API to set the budget on the cluster with the GetCostForecast and NotificationsWithSubscnbers actions
D. Create an AWS Service Catalog portfolio for each tea
E. Add each team's Amazon EMR cluster as an AWS CloudFormation template to their Service Catalog portfolio as a Product
F. Create an Amazon CloudWatch metric for billing Create a custom alert when costs exceed the budgetary threshold.

Answer: BE

NEW QUESTION 120
- (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 onelocation 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 121
- (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 Amazon EC2 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 124
- (Exam Topic 2)
A company's interactive web application uses an Amazon CloudFront distribution to serve images from an Amazon S3 bucket. Occasionally, third-party tools ingest corrupted images into the S3 bucket. This image corruption causes a poor user experience in the application later. The company has successfully implemented and tested Python logic to detect corrupt images.
A solutions architect must recommend a solution to integrate the detection logic with minimal latency between the ingestion and serving.
Which solution will meet these requirements?

A. Use a Lambda@Edge function that is invoked by a viewer-response event.
B. Use a Lambda@Edge function that is invoked by an origin-response event.
C. Use an S3 event notification that invokes an AWS Lambda function.
D. Use an S3 event notification that invokes an AWS Step Functions state machine.

Answer: A

NEW QUESTION 126
- (Exam Topic 2)
A company wants to send data from its on-premises systems to Amazon S3 buckets. The company created the S3 buckets in three different accounts. The company must send the data privately without the data traveling across the internet. The company has no existing dedicated connectivity to AWS
Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

A. Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC
B. Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a public VIF between the on-premises environment and the private VPC
C. Create an Amazon S3 interface endpoint in the networking account
D. Create an Amazon S3 gateway endpoint in the networking account
E. Establish a networking account in the AWS Clou
F. Create a private VPC in the networking account Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account

Answer: AD

NEW QUESTION 131
- (Exam Topic 2)
A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.
The solutions architect created the following IAM policy and attached it to an IAM role:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DownloadUpload",
            "Action": [
                "s3:GetObject",
                "s3:GetObjectVersion",
                "s3:PutObject",
                "s3:PutObjectAcl"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:s3:::BucketName/*"
        },
        {
            "Sid": "KMSAccess",
            "Action": [
                "kms:Decrypt",
                "kms:Encrypt"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:kms:Region:Account:key/Key ID"
        }
    ]
}
```

During tests, me solutions architect was able to successfully get existing test objects m the S3 bucket However, attempts to upload a new object resulted in an error message. The error message stated that me action was forbidden.
Which action must me solutions architect add to the IAM policy to meet all the requirements?

A. Kms:GenerateDataKey
B. KmsGetKeyPolpcy
C. kmsGetPubKKey
D. kms:SKjn

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-access-denied-error-kms/
"An error occurred (AccessDenied) when calling the PutObject operation: Access Denied" This error message indicates that your IAM user or role needs permission for the kms:GenerateDataKey action.

NEW QUESTION 136
- (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{commentId} 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 138
- (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 143
- (Exam Topic 2)
During an audit, a security team discovered that a development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository . The security team wants to automatically find and remediate instances of this security vulnerability
Which solution will ensure that the credentials are appropriately secured automatically?

A. Run a script nightly using AWS Systems Manager Run Command to search for credentials on the development instances If found use AWS Secrets Manager to rotate the credentials.
B. Use a scheduled AWS Lambda function to download and scan the application code from CodeCommit If credentials are found, generate new credentials and store them in AWS KMS
C. Configure Amazon Macie to scan for credentials in CodeCommit repositories If credentials are found, trigger an AWS Lambda function to disable the credentials and notify the user
D. Configure a CodeCommit trigger to invoke an AWS Lambda function to scan new code submissions for credentials If credentials are found, disable them in AWS IAM and notify the user.

Answer: A

NEW QUESTION 148
- (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 multi-Region solution that will 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: C

NEW QUESTION 152
- (Exam Topic 2)
A company with several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsAllActions",
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*"
        },
        {
            "Sid": "DenyCloudTrail",
            "Effect": "Deny",
            "Action": "cloudtrail:*",
            "Resource": "*"
        }
    ]
}
```

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?

A. Add s3:CreateBucket with €Allow€ effect to the SCP.
B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
D. Remove the SCP from account 1111-1111-1111.

Answer: C

Explanation:
 is incorrect - https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
"SCPs are similar to AWS Identity and Access Management (IAM) permission policies and use almost the same syntax. However, an SCP never grants permissions."
SCPs alone are not sufficient to granting permissions to the accounts in your organization. No permissions are granted by an SCP. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

NEW QUESTION 155
- (Exam Topic 2)
A software company is using three AWS accounts for each of its 1 0 development teams The company has developed an AWS CloudFormation standard VPC template that includes three NAT gateways The template is added to each account for each team The company is concerned that network costs will increase each time a new development team is added A solutions architect must maintain the reliability of the company's solutions and minimize operational complexity
What should the solutions architect do to reduce the network costs while meeting these requirements?

A. Create a single VPC with three NAT gateways in a shared services account Configure each account VPC with a default route through a transit gateway to the NAT gateway in the shared services account VPC Remove all NAT gateways from the standard VPC template
B. Create a single VPC with three NAT gateways in a shared services account Configure each account VPC with a default route through a VPC peering connection to the NAT gateway in the shared services account VPC Remove all NAT gateways from the standard VPC template
C. Remove two NAT gateways from the standard VPC template Rely on the NAT gateway SLA to cover reliability for the remaining NAT gateway.
D. Create a single VPC with three NAT gateways in a shared services account Configure a Site-to-Site VPN connection from each account to the shared services account Remove all NAT gateways from the standard VPC template

Answer: A

NEW QUESTION 158
- (Exam Topic 2)
A company has developed APIs that use Amazon API Gateway with Regional endpoints. The APIs call AWS Lambda functions that use API Gateway authentication mechanisms. After a design review, a solutions architect identifies a set of APIs that do not require public access.
The solutions architect must design a solution to make the set of APIs accessible only from a VPC. All APIs need to be called with an authenticated user.
Which solution will meet these requirements with the LEAST amount of effort?

A. Create an internal Application Load Balancer (ALB). Create a target grou
B. Select the Lambda function to cal
C. Use the ALB DNS name to call the API from the VPC.
D. Remove the DNS entry that is associated with the API in API Gatewa
E. Create a hosted zone in Amazon Route 53. Create a CNAME record in the hosted zon
F. Update the API in API Gateway with the CNAME recor
G. Use the CNAME record to call the API from the VPC.
H. Update the API endpoint from Regional to private in API Gatewa
I. Create an interface VPC endpoint in the VP
J. Create a resource policy, and attach it to the AP
K. Use the VPC endpoint to call the API from the VPC.
L. Deploy the Lambda functions inside the VP
M. Provision an EC2 instance, and install an Apache server.From the Apache server, call the Lambda function
N. Use the internal CNAME record of the EC2 instance to call the API from the VPC.

Answer: D

NEW QUESTION 161
- (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 165
- (Exam Topic 2)
A solutions architect uses AWS Organizations to manage several AWS accounts for a company. The full Organizations feature set is activated for the organization. All production AWS accounts exist under an OU that is named "production '' Systems operators have full administrative privileges within these accounts by using IAM roles.
The company wants to ensure that security groups in all production accounts do not allow inbound traffic for TCP port 22. All noncompliant security groups must be remediated immediately, and no new rules that allow port 22 can be created.
Winch solution will meet these requirements?

A. Write an SCP that denies the CreateSecurityGroup action with a condition o(ec2:tngress rule with value 22. Apply the SCP to the 'production' OU.
B. Configure an AWS CloudTrail trail for all accounts Send CloudTrail logs to an Amazon S3 bucket In the Organizations management accoun
C. Configure an AWS Lambda function on the management account with permissions to assume a role in all production accounts to describe and modify security group
D. Configure Amazon S3 to invoke the Lambda function on every PutObject event on the S3 bucket Configure the Lambda function to analyze each CloudTrail event for noncompliant security group actions and to automatically remediate any issues.
E. Create an Amazon EvertBridge (Amazon CloudWatch Events) event bus in the Organizations management accoun
F. Create an AWS Cloud Formation template to deploy configurations that send CreateSecurityGroup events to the even! bus from an production accounts Configure an AWS Lambda function in the management account with permissions to assume a role «i all production accounts to describe and modify security group
G. Configure the event bus to invoke the Lambda function Configure the Lambda function to analyse each event for noncompliant security group actions and to automatically remediate any issues.
H. Create an AWS CloudFormation template to turn on AWS Config Activate the INCOMING_SSH_DISABLED AWS Config managed rule Deploy an AWS Lambda function that will run based on AWS Config findings and will remediate noncompliant resources Deploy the CloudFormation template by using a StackSet that is assigned to the "production" O
I. Apply an SCP to the OU to deny modification of the resources that the CloudFormation template provisions.

Answer: D

NEW QUESTION 168
- (Exam Topic 2)
A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that if a production CloudFormation stack is deleted, important data stored in Amazon RD5 databases or Amazon EBS volumes might also be deleted.
now can the company prevent users from accidentally deleting data m this way?

A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
C. Modify IAM policies to deny deleting RDS and EBS resources that ate lagged with an "aws:cloudformation:stack-name'' tag.
D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Answer: A

NEW QUESTION 169
- (Exam Topic 2)
A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.
The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.
Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.
How can the company solve this problem?

A. Turn on termination protection for the EC2 instances.
B. Update the visibility timeout for the SOS queue to 3 hours.
C. Configure scale-in protection for the instances during processing.
D. Update the redrive policy and set maxReceiveCount to 0.

Answer: D

NEW QUESTION 174
- (Exam Topic 2)

A company has VPC flow logs enabled for its NAT gateway. The company is seeing Action = ACCEPT for inbound traffic that comes from public IP address 198.51.100.2 destined for a private Amazon EC2 instance.

A solutions architect must determine whether the traffic represents unsolicited inbound connections from the internet. The first two octets of the VPC CIDR block are 203.0.

Which set of steps should the solutions architect take to meet these requirements?

A. Open the AWS CloudTrail consol
B. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interfac
C. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
D. Open the Amazon CloudWatch consol
E. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interfac
F. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
G. Open the AWS CloudTrail consol
H. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interfac
I. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
J. Open the Amazon CloudWatch consol
K. Select the log group that contains the NAT gateway's elasticnetwork interface and the private instance's elastic network interfac
L. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

Answer: D

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/vpc-analyze-inbound-traffic-nat-gateway/ by Cloudxie says "select appropriate log"

NEW QUESTION 178
- (Exam Topic 2)
A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS The solution will use Amazon Aurora PostgreSQL. Amazon EC2 Auto Scaling, and Elastic Load Balancing.
Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

A. Enable Aurora Auto Scaling for Aurora Replica
B. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled
C. Enable Aurora Auto Scaling for Aurora writer
D. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled
E. Aurora Auto Scaling for Aurora Replica
F. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
G. Aurora Auto Scaling for Aurora writer
H. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

Explanation:
Aurora Auto Scaling enables your Aurora DB cluster to handle sudden increases in connectivity or workload. When the connectivity or workload decreases, Aurora Auto Scaling removes unnecessary Aurora Replicas so that you don't pay for unused provisioned DB instances

NEW QUESTION 183
- (Exam Topic 2)
A company runs an IoT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MOTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named iot.example.com. The company uses Amazon Route 53 as its DNS service. The company stores the data in Amazon DynamoDB.
On several occasions, the amount of data has overloaded the MOTT broker and has resulted in lost sensor data. The company must improve the reliability of the solution.
Which solution will meet these requirements?

A. Create an Application Load Balancer (ALB) and an Auto Scaling group for the MOTT broke
B. Use the Auto Scaling group as the target for the AL
C. Update the DNS record in Route 53 to an alias recor
D. Point the alias record to the AL
E. Use the MQTT broker to store the data.
F. Set up AWS IoT Core to receive the sensor dat
G. Create and configure a custom domain to connect to AWS IoT Cor
H. Update the DNS record in Route 53 to point to the AWS IoT Core Data-ATS endpoin
I. Configure an AWS IoT rule to store the data.
J. Create a Network Load Balancer (NLB). Set the MQTT broker as the targe
K. Create an AWS Global Accelerator accelerato
L. Set the NLB as the endpoint for the accelerato
M. Update the DNS record in Route 53 to a multivalue answer recor
N. Set the Global Accelerator IP addresses as value
O. Use the MQTT broker to store the data.
P. Set up AWS IoT Greengrass to receive the sensor dat
Q. Update the DNS record in Route 53 to point to the AWS IoT Greengrass endpoin
R. Configure an AWS IoT rule to invoke an AWS Lambda function to store the data.

Answer: C

NEW QUESTION 184
- (Exam Topic 2)
A company runs many workloads on AWS and uses AWS Organizations to manage its accounts. The workloads are hosted on Amazon EC2, AWS Fargate, and AWS Lambda. Some of the workloads have unpredictable demand. Accounts record high usage in some months and low usage in other months.
The company wants to optimize its compute costs over the next 3 years. A solutions architect obtains a 6-month average for each of the accounts across the organization to calculate usage.
Which solution will provide the MOST cost savings for all the organization's compute usage?

A. Purchase Reserved Instances for the organization to match the size and number of the most common EC2 instances from the member accounts.
B. Purchase a Compute Savings Plan for the organization from the management account by using the recommendation at the management account level.
C. Purchase Reserved Instances for each member account that had high EC2 usage according to the data from the last 6 months.
D. Purchase an EC2 Instance Savings Plan for each member account from the management account based on EC2 usage data from the last 6 months.

Answer: A

NEW QUESTION 185
- (Exam Topic 2)
A company has its cloud infrastructure on AWS A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts
What should the solutions architect do to meet these requirements?

A. Use AWS CloudFormation templates Add IAM policies to control the various accounts Deploy the templates across the multiple Regions
B. Use AWS Organizations Deploy AWS CloudFormation templates from the management account Use AWS Control Tower to manage deployments across accounts
C. Use AWS Organizations and AWS CloudFormation StackSets Deploy a CloudFormation template from an account that has the necessary IAM permissions
D. Use nested stacks with AWS CloudFormation templates Change the Region by using nested stacks

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-orga

NEW QUESTION 189
- (Exam Topic 2)
A company is developing a gene reporting device that will collect genomic information to assist researchers with collecting large samples of data from a diverse population. The device will push 8 KB of genomic data every second to a data platform that will need to process and analyze the data and provide information back to researchers The data platform must meet the following requirements:
• Provide near-real-time analytics of the inbound genomic data
• Ensure the data is flexible, parallel, and durable
• Deliver results of processing to a data warehouse
Which strategy should a solutions architect use to meet these requirements?

A. Use Amazon Kinesis Data Firehose to collect the inbound sensor data analyze the data with Kinesis client
B. and save the results to an Amazon RDS instance
C. Use Amazon Kinesis Data Streams to collect the inbound sensor data analyze the data with Kinesis clients and save the results to an Amazon Redshift duster using Amazon EMR
D. Use Amazon S3 to collect the inbound device data analyze the data from Amazon SOS with Kinesis and save the results to an Amazon Redshift duster
E. Use an Amazon API Gateway to put requests into an Amazon SQS queue analyze the data with an AWS Lambda function and save the results » an Amazon Redshift duster using Amazon EMR

Answer: A

NEW QUESTION 191
- (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://1111222223333.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 196
- (Exam Topic 2)
A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times for all task, master, and core nodes The EMR tasks run each morning, starting at 1:00 AM, and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.
The solutions architect must review the architecture and suggest a solution to minimize the compute costs Which solution should the solutions architect

recommend to meet these requirements?

A. Launch all task, master, and core nodes on Spot Instances in an instance flee
B. Terminate the cluster, including all instances, when the processing is completed.
C. Launch the master and core nodes on On-Demand Instance
D. Launch the task nodes on Spot Instances In an instance flee
E. Terminate the cluster, including all instances, when the processing is complete
F. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
G. Continue to launch all nodes on On-Demand Instance
H. Terminate the cluste
I. Including all instances, when the processing Is complete
J. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
K. Launch the master and core nodes on On-Demand Instance
L. Launch the task nodes on Spot Instances In an instance flee
M. Terminate only the task node Instances when the processing is completed Purchase Compute Savings Plans to cover the On-Demand Instance usage.

Answer: B

NEW QUESTION 197
- (Exam Topic 2)
A company is using a lift-and-shift strategy to migrate applications from several on-premises Windows servers to AWS. The Windows servers will be hosted on Amazon EC2 instances in the us-east-1 Region.
The company's security policy allows the installation of migration tools on servers. The migration data must be encrypted in transit and encrypted at rest. The applications are business critical. The company needs to minimize the cutover window and minimize the downtime that results from the migration. The company wants to use Amazon CloudWatch and AWS CloudTrail for monitoring.
Which solution will meet these requirements?

A. Use AWS Application Migration Service (CloudEnsure Migration) to migrate the Windows servers to AW
B. Create a Replication Settings templat
C. Install the AWS Replication Agent on the source servers
D. Use AWS DataSync to migrate the Windows servers to AW
E. Install the DataSync agent on the source server
F. Configure a blueprint for the target server
G. Begin the replication process.
H. Use AWS Server Migration Service (AWS SMS) to migrate the Windows servers to AW
I. Install the SMS Connector on the source server
J. Replicate the source servers to AW
K. Convert the replicated volumes to AMIs to launch EC2 instances.
L. Use AWS Migration Hub to migrate the Windows servers to AW
M. Create a project in Migration Hub.Track the progress of server migration by using the built-in dashboard.

Answer: A

NEW QUESTION 201
- (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 202
- (Exam Topic 2)
A company wants to migrate its data analytics environment from on premises to AWS The environment consists of two simple Node js applications One of the applications collects sensor data and loads it into a MySQL database The other application aggregates the data into reports When the aggregation jobs run. some of the load jobs fail to run correctly
The company must resolve the data loading issue The company also needs the migration to occur without interruptions or changes for the company's customers What should a solutions architect do to meet these requirements?

A. Set up an Amazon Aurora MySQL database as a replication target for the on-premises database Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against theAurora Replica Set up collection endpomts as AWS Lambda functions behind a Network Load Balancer (NLB). and use Amazon RDS Proxy to wnte to the Aurora MySQL database When the databases are synced disable the replication job and restart the Aurora Replica as the primary instanc
B. Point the collector DNS record to the NLB.
C. Set up an Amazon Aurora MySQL database Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora Move the aggregation jobs to run against the Aurora MySQL database Set up collection endpomts behind an Application Load Balancer (ALB) as Amazon EC2 instances in an Auto Scaling group When the databases are synced, point the collector DNS record to the ALB Disable the AWS DMS sync task after the cutover from on premises to AWS
D. Set up an Amazon Aurora MySQL database Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB) and use Amazon RDS Proxy to write to the Aurora MySQL database When the databases are synced, point the collector DNS record to the ALB Disable the AWS DMS sync task after the cutover from on premises to AWS

E. Set up an Amazon Aurora MySQL database Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica Set up collection endpoints as an Amazon Kinesis data stream Use Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database When the databases are synced disable the replication job and restart the Aurora Replica as the primary instance Point the collector DNS record to the Kinesis data stream.

Answer: C

Explanation:
Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.
Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%

NEW QUESTION 205
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

AWS-Certified-Solutions-Architect-Professional Practice Exam Features:

* AWS-Certified-Solutions-Architect-Professional Questions and Answers Updated Frequently

* AWS-Certified-Solutions-Architect-Professional Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-Solutions-Architect-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-Solutions-Architect-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The AWS-Certified-Solutions-Architect-Professional Practice Test Here