Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

    All examinations will be up to date.

* 24/7 Quality Support

    We will provide service round the clock.

* 100% Pass Rate

    Our guarantee that you will pass the exam.

* Unique Gurantee

    If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
- (Exam Topic 1)
A startup company recently migrated a large ecommerce website to AWS. The website has experienced a 70% increase in sales. Software engineers are using a private GitHub repository to manage code. The DevOps learn is using Jenkins for builds and unit testing. The engineers need to receive notifications for bad builds and zero downtime during deployments. The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue.
The software engineers have decided to use AWS CodePipeline to manage their build and deployment process.
Which solution will meet these requirements?

A. Use GitHub websockets to trigger the CodePipeline pipelin
B. Use the Jenkins plugin for AWS CodeBuild to conduct unit testin
C. Send alerts to an Amazon SNS topic for any bad build
D. Deploy in an in-plac
E. all-at-once deployment configuration using AWS CodeDeploy.
F. Use GitHub webhooks to trigger the CodePipeline pipelin
G. Use the Jenkins plugin for AWS CodeBuild to conduct unit testin
H. Send alerts to an Amazon SNS topic for any bad build
I. Deploy in a blue/green deployment using AWS CodeDeploy.
J. Use GitHub websockets to trigger the CodePipeline pipelin
K. Use AWS X-Ray for unit testing and static code analysi
L. Send alerts to an Amazon SNS topic for any bad build
M. Deploy in a blue/green deployment using AWS CodeDeploy.
N. Use GitHub webhooks to trigger the CodePipeline pipelin
O. Use AWS X-Ray for unit testing and static code analysi
P. Send alerts to an Amazon SNS topic for any bad build
Q. Deploy in an in-place, all-at-once deployment configuration using AWS CodeDeploy.

Answer: B

NEW QUESTION 2
- (Exam Topic 1)
A company hosts a large on-premises MySQL database at its main office that supports an issue tracking system used by employees around the world. The company already uses AWS for some workloads and has created an Amazon Route 53 entry for the database endpoint that points to the on-premises database. Management is concerned about the database being a single point of failure and wants a solutions architect to migrate the database to AWS without any data loss or downtime.
Which set of actions should the solutions architect implement?

A. Create an Amazon Aurora DB cluste
B. Use AWS Database Migration Service (AWS DMS) to do a full load from the on-premises database lo Auror
C. Update the Route 53 entry for the database to point to the Aurora cluster endpoin
D. and shut down the on-premises database.
E. During nonbusiness hours, shut down the on-premises database and create a backu
F. Restore this backup to an Amazon Aurora DB cluste
G. When the restoration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
H. Create an Amazon Aurora DB cluste
I. Use AWS Database Migration Service (AWS DMS) to do a full load with continuous replication from the on-premises database to Auror
J. When the migration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
K. Create a backup of the database and restore it to an Amazon Aurora multi-master cluste
L. This Aurora cluster will be in a master-master replication configuration with the on-premises databas
M. Update the Route 53 entry for the database to point to the Aurora cluster endpoin
N. and shut down the on-premises database.

Answer: C

Explanation:
"Around the world" eliminates possibility for the maintenance window at night. The other difference is ability to leverage continuous replication in MySQL to Aurora case.

NEW QUESTION 3
- (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 4
- (Exam Topic 1)
A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.
The company has the following DNS resolution requirements:
• On-premises systems should be able to resolve and connect to cloud.example.com.
• All VPCs should be able to resolve cloud.example.com.
There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway. Which architecture should the company use to meet these requirements with the HIGHEST performance?

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

Answer: D

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

NEW QUESTION 5
- (Exam Topic 1)
An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP. Java, or Ruby web applications, are no longer actively developed, and serve little traffic.
Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs?

A. Deploy the applications lo single-instance AWS Elastic Beanstalk environments without a load balancer.
B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.
C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.
D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

Answer: C

NEW QUESTION 6
- (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 thespecific AWS Organizations OU that will use the shared netwo
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 netwo
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 7
- (Exam Topic 1)
A development team has created a new flight tracker application that provides near-real-time data to users. The application has a front end that consists of an Application Load Balancer (ALB) in front of two large Amazon EC2 instances in a single Availability Zone. Data is stored in a single Amazon RDS MySQL DB instance. An Amazon Route 53 DNS record points to the ALB.
Management wants the development team to improve the solution to achieve maximum reliability with the least amount of operational overhead.
Which set of actions should the team take?

A. Create RDS MySQL read replica
B. Deploy the application to multiple AWS Region
C. Use a Route 53 latency-based routing policy to route to the application.

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

Answer: D

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

NEW QUESTION 8
- (Exam Topic 1)
A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.
The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.
Which storage strategy is the MOST cost-effective and meets the design requirements?

A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieva
B. Configure a lifecycle policy to delete data older than 120 days.
C. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scal
D. Configure the DynamoOB Time to Live (TTL) feature to delete records older than 120 days.
E. Design the application to store each incoming record in a single table in an Amazon RDS MySQL databas
F. Run a nightly cron job that executes a query to delete any records older than 120 days.
G. Design the application to batch incoming records before writing them to an Amazon S3 bucke
H. Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadatasearch feature to retrieve the dat
I. Configure a lifecycle policy to delete the data after 120 days.

Answer: B

Explanation:
DynamoDB with TTL, cheaper for sustained throughput of small items + suited for fast retrievals. S3 cheaper for storage only, much higher costs with writes. RDS not designed for this use case.

NEW QUESTION 9
- (Exam Topic 1)
A company has developed an application that is running Windows Server on VMware vSphere VMs that the company hosts or premises. The application data is stored in a proprietary format that must be read through the application. The company manually provisioned the servers and the application.
As pan of us disaster recovery plan, the company warns the ability to host its application on AWS temporarily me company's on-premises environment becomes unavailable The company wants the application to return to on-premises hosting after a disaster recovery event is complete The RPO 15 5 minutes.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Configure AWS DataSyn
B. Replicate the data lo Amazon Elastic Block Store (Amazon EBS) volumes When the on-premises environment is unavailable, use AWS CloudFormation templates to provision Amazon EC2 instances and attach the EBS volumes
C. Configure CloudEndure Disaster Recovery Replicate the data to replication Amazon EC2 instances that are attached to Amazon Elastic Block Store (Amazon EBS) volumes When the on-premises environment is unavailable, use CloudEndure to launch EC2 instances that use the replicated volumes.
D. Provision an AWS Storage Gateway We gatewa
E. Recreate the data lo an Amazon S3 bucke
F. When the on-premises environment is unavailable, use AWS Backup to restore the data to Amazon Elastic Block Store (Amazon EBS) volumes and launch Amazon EC2 instances from these EBS volumes
G. Provision an Amazon FS* for Windows File Server file system on AWS Replicate :ne data to the «e system When the on-premoes environment is unavailable, use AWS CloudFormation templates to provision Amazon EC2 instances and use AWS :CloudFofmation::lnit commands to mount the Amazon FSx file shares

Answer: D

NEW QUESTION 10
- (Exam Topic 1)
A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.
Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts
Which solution meets these requirements with the LEAST amount of operational overhead?

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

in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: B

NEW QUESTION 10
- (Exam Topic 1)
A group of research institutions and hospitals are in a partnership to study 2 PBs of genomic data. The institute that owns the data stores it in an Amazon S3 bucket and updates it regularly. The institute would like to give all of the organizations in the partnership read access to the data. All members of the partnership are extremely cost-conscious, and the institute that owns the account with the S3 bucket is concerned about covering the costs tor requests and data transfers from Amazon S3.
Which solution allows for secure datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers'?

A. Ensure that all organizations in the partnership have AWS account
B. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the dat
C. Have the organizations assume and use that read role when accessing the data.
D. Ensure that all organizations in the partnership have AWS account
E. Create a bucket policy on the bucket that owns the data The policy should allow the accounts in the partnership read access to the bucke
F. Enable Requester Pays on the bucke
G. Have the organizations use their AWS credentials when accessing the data.
H. Ensure that all organizations in the partnership have AWS account
I. Configure buckets in each of the accounts with a bucket policy that allows the institute that owns the data the ability to write to the bucket Periodically sync the data from the institute's account to the other organization
J. Have the organizations use their AWS credentials when accessing the data using their accounts
K. Ensure that all organizations in the partnership have AWS account
L. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the dat
M. Enable Requester Pays on the bucke
N. Have the organizations assume and use that read role when accessing the data.

Answer: B

Explanation:
In general, bucket owners pay for all Amazon S3 storage and data transfer costs associated with their bucket. A bucket owner, however, can configure a bucket to be a Requester Pays bucket. With Requester Pays buckets, the requester instead of the bucket owner pays the cost of the request and the data download from the bucket. The bucket owner always pays the cost of storing data. If you enable Requester Pays on a bucket, anonymous access to that bucket is not allowed.
https://docs.aws.amazon.com/AmazonS3/latest/userguide/RequesterPaysExamples.html

NEW QUESTION 15
- (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 19
- (Exam Topic 1)
A solution architect is designing an AWS account structure for a company that consists of multiple terms. All the team will work in the same AWS Region. The company needs a VPC that is connected to the on-premises network. The company expects less than 50 Mbps of total to and from the on-premises network.
Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO)

A. Create an AWS CloudFormation template that provisions a VPC and the required subnet
B. Deploy the template to each AWS account
C. Create an AWS CloudFormabon template that provisions a VPC and the required subnet
D. Deploy the template to a shared services accoun
E. Share the subnets by using AWS Resource Access Manager
F. Use AWS Transit Gateway along with an AWS Site-to-Site VPN for connectivity to the on-premises networ
G. Share the transit gateway by using AWS Resource Access Manager
H. Use AWS Site-to-Site VPN for connectivity to the on-premises network
I. Use AWS Direct Connect for connectivity to the on-premises network.

Answer: BD

NEW QUESTION 24

- (Exam Topic 1)

A finance company is running its business-critical application on current-generation Linux EC2 instances The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.

Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

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

Answer: B

Explanation:

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

NEW QUESTION 27

- (Exam Topic 1)

A company is running a web application on Amazon EC2 instances in a production AWS account. The company requires all logs generated from the web application to be copied to a central AWS account (or analysis and archiving. The company's AWS accounts are currently managed independently. Logging agents are configured on the EC2 instances to upload the tog files to an Amazon S3 bucket in the central AWS account.

A solutions architect needs to provide access for a solution that will allow the production account to store log files in the central account. The central account also needs to have read access to the tog files.

What should the solutions architect do to meet these requirements?

A. Create a cross-account role in the central accoun
B. Assume the role from the production account when the logs are being copied.
C. Create a policy on the S3 bucket with the production account ID as the principa
D. Allow S3 access from a delegated user.
E. Create a policy on the S3 bucket with access from only the CIDR range of the EC2 instances in the production accoun
F. Use the production account ID as the principal.
G. Create a cross-account role in the production accoun
H. Assume the role from the production account when the logs are being copied.

Answer: B

NEW QUESTION 29

- (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 32

- (Exam Topic 1)

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

Which of the following is the MOST reliable approach to meet the requirements?

A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.

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

Answer: B

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

NEW QUESTION 33
- (Exam Topic 1)
A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS lor PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.
What changes to the current architecture will reduce operational overhead and support the product release?

A. Create an EC2 Auto Scaling group behind an Application Load Balance
B. Create additional read replicas for the DB instanc
C. Create Amazon Kinesis data streams and configure the application services lo use the data stream
D. Store and serve static content directly from Amazon S3.
E. Create an EC2 Auto Scaling group behind an Application Load Balance
F. Deploy the DB instance in Multi-AZ mode and enable storage auto scalin
G. Create Amazon Kinesis data streams and configure the application services to use the data stream
H. Store and serve static content directly from Amazon S3.
I. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balance
J. Deploy the DB instance in Multi-AZ mode and enable storage auto scalin
K. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluste
L. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
M. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balance
N. Create additional read replicas for the DB instanc
O. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluste
P. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

Answer: D

Explanation:
Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

NEW QUESTION 35
- (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 an Amazon 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 39
- (Exam Topic 1)

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

A. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon MQ queue.
B. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue.
C. Invoke an AWS Lambda function to perform image processing when a message is available in the queue.
D. Invoke an S3 Batch Operations job to perform image processing when a message is available in the queue.
E. Send a push notification to the mobile app by using Amazon Simple Notification Service (AmazonSNS) when processing is complete.
F. Send a push notification to the mobile app by using Amazon Simple Email Service (Amazon SES) when processing is complete.

Answer: BCE

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

NEW QUESTION 42
- (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 44
- (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 47
- (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 51
- (Exam Topic 1)
A solutions architect is responsible (or redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.

Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of (he system while decreasing the processing latency and minimizing costs?

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

Answer: B

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

NEW QUESTION 55
- (Exam Topic 1)
A team collects and routes behavioral data for an entire company. The company runs a Multi-AZ VPC environment with public subnets, private subnets, and in internet gateway Each public subnet also contains a NAT gateway Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads run in private subnets.

A solutions architect must review the infrastructure The solutions architect needs to reduce costs and maintain the function of the applications. The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.

What should the solutions architect do to meet these requirements?

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

Answer: D

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

NEW QUESTION 56
- (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 60
- (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 63
- (Exam Topic 1)
A company is running an application distributed over several Amazon EC2 instances in an Auto Seating 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 log
B. Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.
C. 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.
D. 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
E. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the sourc
F. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elasticsearch 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 https://docs.aws.amazon.com/vpc/latest/mirroring/what-is-traffic-mirroring.html

NEW QUESTION 64
- (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 68
- (Exam Topic 1)
A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.
After initial deployment, the company observes 1.000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.
Which approach should the company take to secure its API?

A. Create an Amazon CloudFront distribution with the API as the origi
B. Create an AWS WAF web ACL with a rule to block clients "hat submit more than five requests per da

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

Answer: D

Explanation:
"A usage plan specifies who can access one or more deployed API stages and methods—and also how much and how fast they can access them. The plan uses API keys to identify API clients and meters access to the associated API stages for each key. It also lets you configure throttling limits and quota limits that are enforced on individual client API keys."
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html

NEW QUESTION 73
- (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 74
- (Exam Topic 1)
A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an
on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.
What should the solutions architect do to meet these requirements?

A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows File Server Use the SMB share to host the VMware data stor
B. Use VM Import/Export to move the VMs to Amazon EC2.
C. Use the VMware vSphere client to export the application as an image in Open Virealization Format (OVF) format Create an Amazon S3 bucket to store the image in the destination AWS Regio
D. Create and apply an IAM role for VM Import Use the AWS CLI to run the EC2 import command.
E. Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFSJ shar
F. Create a backup copy to the shared folde
G. Sign in to the AWS Management Console and create an AMI from the backup copy Launch an EC2 instance that is based on the AMI.
H. Create a managed-instance activation for a hybrid environment in AWS Systems Manage
I. Download and install Systems Manager Agent on the on-premises VM Register the VM with Systems Manager to be a managed instance Use AWS Backup to create a snapshot of the VM and create an AM
J. Launch an EC2 instance that is based on the AMI

Answer: B

Explanation:
https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html
- Export an OVF Template

- Create / use an Amazon S3 bucket for storing the exported images. The bucket must be in the Region where you want to import your VMs.
- Create an IAM role named vmimport.
- You'll use AWS CLI to run the import commands. https://aws.amazon.com/premiumsupport/knowledge-center/import-instances/

NEW QUESTION 77
- (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 80
- (Exam Topic 1)
A company is moving a business-critical multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that frequently fails to maintain the application uptime SLA of 99.95%. A solutions architect must re-architect the application to ensure that it can meet or exceed the SLA.
The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.
Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?

A. Migrate the database to a PostgreSQL database in Amazon EC2. Host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
B. Allocate an Amazon Workspaces Workspace for each end user to improve the user experience.
C. Migrate the database to an Amazon RDS Aurora PostgreSQL configuratio
D. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balance
E. Use Amazon AppStream 2.0 to improve the user experience.
F. Migrate the database to an Amazon RDS PostgreSQL Mulli-AZ configuratio
G. Host the application and presentation layers in automatically scaled AWS Fargate containers behind a Network Load Balance
H. Use Amazon ElastiCache to improve the user experience.
I. Migrate the database to an Amazon Redshift cluster with at least two node
J. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
K. Use Amazon CloudFront to improve the user experience.

Answer: B

Explanation:
Aurora would improve availability that can replicate to multiple AZ (6 copies). Auto scaling would improve the performance together with a ALB. AppStream is like Citrix that deliver hosted Apps to users.

NEW QUESTION 85
- (Exam Topic 1)
A North American company with headquarters on the East Coast is deploying a new web application running on Amazon EC2 in the us-east-1 Region. The application should dynamically scale to meet user demand and maintain resiliency. Additionally, the application must have disaster recovery capabilities in an active-passive configuration with the us-west-1 Region.
Which steps should a solutions architect take after creating a VPC in the us-east-1 Region?

A. Create a VPC in the us-west-1 Regio
B. Use inter-Region VPC peering to connect both VPC
C. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in theus-east-1 Regio
D. Deploy EC2 instances across multiple AZs in each Region as part of an Auto Scaling group spanning both VPCs and served by the ALB.
E. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC inthe us-east-1 Regio
F. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the AL
G. Deploy the same solution to the us-west-1 Region Create an Amazon Route 53 record set with a failover routing policy and health checks enabled to provide high availability across both Regions.
H. Create a VPC in the us-west-1 Regio
I. Use inter-Region VPC peering to connect both VPCs Deploy an Application Load Balancer (ALB) that spans both VPCs Deploy EC2 instances across multiple Availability Zones as part of an Auto Scaling group in each VPC served by the AL
J. Create an Amazon Route 53 record that points to the ALB.
K. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Regio
L. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the AL
M. Deploy the same solution to the us-west-1 Regio
N. Create separate Amazon Route 53 records in each Region that point to the ALB in the Regio
O. Use Route 53 health checks to provide high availability across both Regions.

Answer: B

Explanation:
A new web application in a active-passive DR mode. a Route 53 record set with a failover routing policy.

NEW QUESTION 88
- (Exam Topic 1)
A company needs to run a software package that has a license that must be run on the same physical host for the duration of Its use. The software package is only going to be used for 90 days The company requires patching and restarting of all instances every 30 days
How can these requirements be met using AWS?

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

Answer: B

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

NEW QUESTION 92
- (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 95
- (Exam Topic 1)
A company requires that all internal application connectivity use private IP addresses. To facilitate this policy, a solutions architect has created interface endpoints to connect to AWS public services. Upon testing, the solutions architect notices that the service names are resolving to public IP addresses, and that internal services cannot connect to the interface endpoints.
Which step should the solutions architect take to resolve this issue?

A. Update the subnet route table with a route to the interface endpoint.
B. Enable the private DNS option on the VPC attributes.
C. Configure the security group on the interface endpoint to allow connectivity to the AWS services.
D. Configure an Amazon Route 53 private hosted zone with a conditional forwarder for the internal application.

Answer: C

Explanation:
https://docs.aws.amazon.com/vpc/latest/privatelink/vpce-interface.html

NEW QUESTION 96
- (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 97
- (Exam Topic 1)
A company wants to control its cost of Amazon Athena usage The company has allocated a specific monthly budget for Athena usage A solutions architect must design a solution that will prevent the company from exceeding the budgeted amount
Which solution will moot these requirements?

A. Use AWS Budget
B. Create an alarm (or when the cost of Athena usage reaches the budgeted amount for the mont
C. Configure AWS Budgets actions to deactivate Athena until the end of the month.
D. Use Cost Explorer to create an alert for when the cost of Athena usage reaches the budgeted amount for the mont
E. Configure Cost Explorer to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic.
F. Use AWS Trusted Advisor to track the cost of Athena usag
G. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to deactivate Athena until the end of the month whenever the cost reaches the budgeted amount for the month
H. Use Athena workgroups to set a limit on the amount of data that can be scanne
I. Set a limit that is appropriate for the monthly budget and the current pricing for Athena.

Answer: D

NEW QUESTION 98
- (Exam Topic 1)
A developer reports receiving an Error 403: Access Denied message when they try to download an object from an Amazon S3 bucket. The S3 bucket is accessed using an S3 endpoint inside a VPC. and is encrypted with an AWS KMS key. A solutions architect has verified that (he developer is assuming the correct IAM role in the account that allows the object to be downloaded. The S3 bucket policy and the NACL are also valid.
Which additional step should the solutions architect take to troubleshoot this issue?

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

Answer: B

NEW QUESTION 99
- (Exam Topic 1)
A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalogue page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.
Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times.
Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Select TWO.)

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

Answer: BE

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

NEW QUESTION 104
- (Exam Topic 1)
A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.
The company wants to stop using email in the application and move to a serverless application model. Which architecture solution meets these requirements?

A. Use AWS Batch to configure the different tasks required lo ship a packag
B. Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping labe
C. Once that label is scanne

D. as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.B.
E. When a new order is created, store the order information in Amazon SQ
F. Have AWS Lambda check the queue every 5 minutes and process any needed wor
G. When an order needs to be shipped, have Lambda print the label in the warehous
H. Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon SOS.
I. Update the application to store new order information in Amazon DynamoD
J. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress," and print a package label to the warehous
K. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.
L. Store new order information in Amazon EF
M. Have instances pull the new information from the NFS and send that information to printers in the warehous
N. Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.

Answer: C

NEW QUESTION 109
- (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 111
- (Exam Topic 1)
A company is running an application on Amazon EC2 instances in three environments; development, testing, and production. The company uses AMIs to deploy the EC2 instances. The company builds the AMIs by using custom deployment scripts and infrastructure orchestration tools for each release in each environment. The company is receiving errors in its deployment process. Errors appear during operating system package downloads and during application code installation from a third-party Git hosting service. The company needs deployments to become more reliable across all environments.
Which combination of steps will meet these requirements? (Select THREE).

A. Mirror the application code to an AWS CodeCommit Git repositor
B. Use the repository to build EC2 AMIs.
C. Produce multiple EC2 AMI
D. one for each environment, for each release.
E. Produce one EC2 AMI for each release for use across all environments.
F. Mirror the application code to a third-party Git repository that uses Amazon S3 storag
G. Use therepository for deployment.
H. Replace the custom scripts and tools with AWS CodeBuil
I. Update the infrastructure deployment process to use EC2 Image Builder.

Answer: ACE

NEW QUESTION 113
- (Exam Topic 1)
A company is migrating applications from on premises to the AWS Cloud. These applications power the company's internal web forms. These web forms collect data for specific events several times each quarter. The web forms use simple SQL statements to save the data to a local relational database.
Data collection occurs for each event, and the on-premises servers are idle most of the time. The company needs to minimize the amount of idle infrastructure that supports the web forms.
Which solution will meet these requirements?

A. Use Amazon EC2 Image Builder to create AMIs for the legacy server
B. Use the AMIs to provision EC2 instances to recreate the applications in the AWS.Clou
C. Place an Application Load Balancer (ALB) in front of the EC2 instance
D. Use Amazon Route 53 to point the DNS names of the web forms to the ALB.
E. Create one Amazon DynamoDB table to store data for all the data input Use the application form name as the table key to distinguish data item
F. Create an Amazon Kinesis data stream to receive the data input and store the input in DynamoD
G. Use Amazon Route 53 to point the DNS names of the web forms to the Kinesis data stream's endpoint.
H. Create Docker images for each server of the legacy web form application
I. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargat
J. Place an Application Load Balancer in front of the ECS cluste
K. Use Fargate task storage to store the web form data.
L. Provision an Amazon Aurora Serverless cluste
M. Build multiple schemas for each web form's data storag
N. Use Amazon API Gateway and an AWS Lambda function to recreate the data input form
O. Use Amazon Route 53 to point the DNS names of the web forms to their corresponding API Gateway endpoint.

Answer: D

Explanation:
Provision an Amazon Aurora Serverless cluster. Build multiple schemas for each web forms data storage. Use Amazon API Gateway and an AWS Lambda function to recreate the data input forms. Use Amazon Route 53 to point the DNS names of the web forms to their corresponding API Gateway endpoint.

NEW QUESTION 118
- (Exam Topic 1)
A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.
Which set of actions should a solutions architect take to meet these requirements?

A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instance
B. Use Systems Manager to generate patch compliance reports.
C. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
D. Use Amazon OuickSight integration with OpsWorks to generate patch compliance reports.
E. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to apply patches by scheduling an AWS Systems Manager patch remediation jo
F. Use Amazon Inspector to generate patch compliance reports.
G. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
H. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

Answer: A

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

NEW QUESTION 122
- (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 123
- (Exam Topic 2)
A company that runs applications on AWS recently subscribed to a new software-as-a-service (SaaS) data vendor. The vendor provides the data by way of a REST API that the vendor hosts in its AWS environment The vendor offers multiple options for connectivity to the API and Is working with the company to find the best way to connect.
The company's AWS account does not allow outbound internet access from Its AWS environment The vendor's services run on AWS in the same AWS Region as the company's applications
A solutions architect must Implement connectivity to the vendor's API so that the API is highly available In the company's VPC.
Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 127
- (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

NEW QUESTION 132

- (Exam Topic 2)
A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises intra structure that consists of physical machines and VMs that host numerous applications.

The company must capture details about the system configuration. system performance. running processure and network coi.net lions of its o. -premises ,on boards. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.

Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs.
B. Assess the existing applications by installing AWS Systems Manager Agent on the physical machines and VMs
C. Group servers into applications for migration by using AWS Systems Manager Application Manager.
D. Group servers into applications for migration by using AWS Migration Hub.
E. Generate recommended instance types and associated costs by using AWS Migration Hub.
F. Import data about server sizes into AWS Trusted Adviso
G. Follow the recommendations for cost optimization.

Answer: BDF

NEW QUESTION 137
- (Exam Topic 2)
A company is running a three-tier web application in an on-premises data center. The frontend is served by an Apache web server, the middle tier is a monolithic Java application, and the storage tier is a PostgreSOL database.

During a recent marketing promotion, customers could not place orders through the application because the application crashed An analysis showed that all three tiers were overloaded. The application became unresponsive, and the database reached its capacity limit because of read operations. The company already has several similar promotions scheduled in the near future.

A solutions architect must develop a plan for migration to AWS to resolve these issues. The solution must maximize scalability and must minimize operational effort.

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

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

Answer: BCF

NEW QUESTION 142
- (Exam Topic 2)
A fleet of Amazon ECS instances is used to poll an Amazon SQS queue and update items in an Amazon DynamoDB database Items in the table are not being updated, and the SQS queue Is filling up Amazon CloudWatch Logs are showing consistent 400 errors when attempting to update the table The provisioned write capacity units are appropriately configured, and no throttling is occurring
What is the LIKELY cause of the failure*?

A. The ECS service was deleted
B. The ECS configuration does not contain an Auto Scaling group
C. The ECS instance task execution IAM role was modified
D. The ECS task role was modified

Answer: D

NEW QUESTION 145
- (Exam Topic 2)
A company has a new security policy. The policy requires the company to log any event that retrieves data from Amazon S3 buckets. The company must save these audit logs in a dedicated S3 bucket. The company created the audit logs S3 bucket in an AWS account that is designated for centralized logging. The S3 bucket has a bucket policy that allows write-only cross-account access A solutions architect must ensure that all S3 object-level access is being logged for current S3 buckets and future S3 buckets. Which solution will meet these requirements?

A. Enable server access logging for all current S3 bucket
B. Use the audit logs S3 bucket as a destination foraudit logs
C. Enable replication between all current S3 buckets and the audit logs S3 bucket Enable S3 Versioning in the audit logs S3 bucket
D. Configure S3 Event Notifications for all current S3 buckets to invoke an AWS Lambda function every time objects are accessed . Store Lambda logs in the audit logs S3 bucket.
E. Enable AWS CloudTrai
F. and use the audit logs S3 bucket to store logs Enable data event logging for S3 event sources, current S3 buckets, and future S3 buckets.

Answer: D

NEW QUESTION 146
- (Exam Topic 2)
A company uses AWS Organizations with a single OU named Production to manage multiple accounts All accounts are members of the Production OU Administrators use deny list SCPs in the root of the organization to manage access to restricted services.

The company recently acquired a new business unit and invited the new unit's existing AWS account to the
organization Once onboarded the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.
Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

A. Remove the organization's root SCPs that limit access to AWS Config Create AWS Service Catalog products for the company's standard AWS Config rules and deploy them throughout the organization, including the new account.
B. Create a temporary OU named Onboarding for the new account Apply an SCP to the Onboarding OU to allow AWS Config actions Move the new account to the Production OU when adjustments to AWS Config are complete
C. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only Temporarily apply an SCP to the organization's root that allows AWS Config actions for principals only in the new account.
D. Create a temporary OU named Onboarding for the new account Apply an SCP to the Onboarding OU to allow AWS Config action
E. Move the organization's root SCP to the Production O
F. Move the new account to the Production OU when adjustments to AWS Config are complete.

Answer: D

NEW QUESTION 151
- (Exam Topic 2)
A company is planning to migrate an application from on premises to the AWS Cloud. The company will begin the migration by moving the application's underlying data storage to AWS The application data is stored on a shared tie system on premises, and the application servers connect to the shared We system through SMB.
A solutions architect must implement a solution that uses an Amazon S3 bucket tor shared storage Until the application Is fully migrated and code is rewritten to use native Amazon S3 APIs, the application must continue to have access to the data through SMB The solutions architect must migrate the application data to AWS to its new location while still allowing the on-premises application to access the data.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 152
- (Exam Topic 2)
A company has deployed an application to multiple environments in AWS. including production and testing the company has separate accounts for production and testing, and users are allowed to create additional
application users for team members or services. as needed. The security team has asked the operations team tor better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments
Which of the following options would MOST securely accomplish this goal?

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

Answer: A

NEW QUESTION 156
- (Exam Topic 2)
A retail company is running an application that stores invoice files in an Amazon S3 bucket and metadata about the files in an Amazon DynamoDB table. The application software runs in both us-east-1 and eu-west-1 The S3 bucket and DynamoDB table are in us-east-1. The company wants to protect itself from data corruption and loss of connectivity to either Region
Which option meets these requirements?

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

Answer: B

NEW QUESTION 158
- (Exam Topic 2)
A Solutions Architect is constructing a containerized.NET Core application for AWS Fargate. The application's backend needs a high-availability version of

Microsoft SQL Server. All application levels must be extremely accessible. The credentials associated with the SQL Server connection string should not be saved to disk inside the.NET Core front-end containers.
Which tactics should the Solutions Architect use to achieve these objectives?

A. Set up SQL Server to run in Fargate with Service Auto Scalin
B. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server running in Fargat
C. Specify the ARN of the secret in AWS Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection strin
D. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
E. Create a Multi-AZ deployment of SQL Server on Amazon RD
F. Create a secret in AWS Secrets Manager for the credentials to the RDS databas
G. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manage
H. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection strin
I. Set up the .NET Core service in Fargate using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
J. Create an Auto Scaling group to run SQL Server on Amazon EC2. Create a secret in AWS Secrets Manager for the credentials to SQL Server running on EC2. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server on EC2. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection strin
K. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
L. Create a Multi-AZ deployment of SQL Server on Amazon RD
M. Create a secret in AWS Secrets Manager for the credentials to the RDS databas
N. Create non- persistent empty storage for the .NET Core containers in the Fargate task definition to store the sensitive informatio
O. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manage
P. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be written to the non-persistent empty storage on startup for reading into the application to construct the connection strin
Q. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.

Answer: B

Explanation:
Secrets Manager natively supports SQL Server on RDS. No real need to create additional 'ephemeral storage' to fetch credentials, as these can be injected to containers as environment variables. https://aws.amazon.com/premiumsupport/knowledge-center/ecs-data-security-container-task/

NEW QUESTION 161
- (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 166
- (Exam Topic 2)
A solutions architect has been assigned to migrate a 50 TB Oracle data warehouse that contains sales data from on-premises to Amazon Redshift Major updates to the sales data occur on the final calendar day of the month For the remainder of the month, the data warehouse only receives minor daily updates and is primarily used for reading and reporting Because of this the migration process must start on the first day of the month and must be complete before the next set of updates occur. This provides approximately 30 days to complete the migration and ensure that the minor daily changes have been synchronized with the Amazon Redshift data warehouse Because the migration cannot impact normal business network operations, the bandwidth allocated to the migration for moving data over the internet is 50 Mbps The company wants to keep data migration costs low
Which steps will allow the solutions architect to perform the migration within the specified timeline?

A. Install Oracle database software on an Amazon EC2 instance Configure VPN connectivity between AWS and the company's data center Configure the Oracle database running on Amazon EC2 to join the Oracle Real Application Clusters (RAC) When the Oracle database on Amazon EC2 finishes synchronizing, create an AWS DMS ongoing replication task to migrate the data from the Oracle database on Amazon EC2 to Amazon Redshift Verify the data migration is complete and perform the cut over to Amazon Redshift.
B. Create an AWS Snowball import job Export a backup of the Oracle data warehouse Copy the exported data to the Snowball device Return the Snowball device to AWS Create an Amazon RDS for Oracle database and restore the backup file to that RDS instance Create an AWS DMS task to migrate the data from the RDS for Oracle database to Amazon Redshift Copy daily incremental backups from Oracle in the data center to the RDS for Oracle database over the internet Verify the

data migration is complete and perform the cut over to Amazon Redshift.
C. Install Oracle database software on an Amazon EC2 instance To minimize the migration time configure VPN connectivity between AWS and the company's data center by provisioning a 1 Gbps AWS Direct Connect connection Configure the Oracle database running on Amazon EC2 to be a read replica of the data center Oracle database Start the synchronization process between the company's on-premises data center and the Oracle database on Amazon EC2 When the Oracle database on Amazon EC2 is synchronized with the on-premises database create an AWS DMS ongoing replication task from the Oracle database read replica that is running on Amazon EC2 to Amazon Redshift Verify the data migration is complete and perform the cut over to Amazon Redshift.
D. Create an AWS Snowball import jo
E. Configure a server in the company€™s data center with an extraction agen
F. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schem
G. Create a new project in AWS SCT using the registered data extraction agen
H. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing change
I. Copy data to the Snowball device and return the Snowball device to AW
J. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshif
K. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

Answer: D

Explanation:
Create an AWS Snowball import job. Configure a server in the company€™s data center with an extraction agent. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema. Create a new project in AWS SCT using the registered data extraction agent. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes. Copy data to the Snowball device and return the Snowball device to AWS. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift. Verify that the data migration is complete and perform the cut over to Amazon Redshift.
https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/

NEW QUESTION 168
- (Exam Topic 2)
A company is running an application in the AWS Cloud. The application runs on containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost. Which solution will meet these requirements with the LEAST amount of operational overhead?

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

Answer: B

NEW QUESTION 171
- (Exam Topic 2)
A company has an on-premises Microsoft SQL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS. Which solution meets these requirements Most cost effectively?

A. Create a new S3 bucket Deploy an AWS Storage Gateway file gateway within the VPC that is connected to the Direct Connect connectio
B. Create a new SMB file shar
C. Write nightly database exports to the new SMB file share.
D. Create an Amzon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the Direct Connect connectio
E. Create a new SMB file shar
F. Write nightly database exports to an SMB file share on the Amazon FSx file system Enable backups.
G. Create an Amazon FSx for Windows File Server Multi-AZ system within the VPC that is connected to the Direct Connect connectio
H. Create a new SMB file shar
I. Write nightly database exports to an SMB file share on the Amazon FSx file syste
J. Enable nightly backups.
K. Create a new S3 bucket
L. Deploy an AWS Storage Gateway volume gateway within the VPC that is connected to the Direct Connect connectio
M. Create a new SMB file shar
N. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.

Answer: A

NEW QUESTION 176
- (Exam Topic 2)
A company is running a workload that consists of thousands of Amazon EC2 instances The workload is running in a VPC that contains several public subnets and private subnets The public subnets have a route for 0 0 0 0/0 to an existing internet gateway. The private subnets have a route for 0 0 0 0/0 to an existing NAT gateway
A solutions architect needs to migrate the entire fleet of EC2 instances to use IPv6 The EC2 instances that are in private subnets must not be accessible from the public internet
What should the solutions architect do to meet these requirements?

A. Update the existing VPC and associate a custom IPv6 CIDR block with the VPC and all subnets Update all the VPC route tables and add a route for /0 to the internet gateway
B. Update the existing VP
C. and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets Update the VPC route tables for all private subnets, and add a route for /0 to the NAT gateway
D. Update the existing VP
E. and associate an Amazon-provided IPv6 CIDR block with the VPC and ail subnets Create an egress-only internet gateway Update the VPC route tables for all private subnets, and add a route for /0 to the egress-only internet gateway
F. Update the existing VPC and associate a custom IPv6 CIDR block with the VPC and all subnets Create a new NAT gateway, and enable IPv6 support Update the VPC route tables for all private subnets and add a route for 70 to the IPv6-enabled NAT gateway.

Answer: C

NEW QUESTION 177
- (Exam Topic 2)
A company has multiple business units Each business unit has its own AWS account and runs a single website within that account. The company also has a single logging account. Logs from each business unit website are aggregated into a single Amazon S3 bucket in the logging account. The S3 bucket policy provides each business unit with access to write data into the bucket and requires data to be encrypted.
The company needs to encrypt logs uploaded into the bucket using a Single AWS Key Management Service
{AWS KMS) CMK The CMK that protects the data must be rotated once every 365 days
Which strategy is the MOST operationally efficient for the company to use to meet these requirements?

A. Create a customer managed CMK ri the logging account Update the CMK key policy to provide access to the logging account only Manually rotate the CMK every 365 days.
B. Create a customer managed CMK in the logging accoun
C. Update the CMK key policy to provide access to the logging account and business unit account
D. Enable automatic rotation of the CMK
E. Use an AWS managed CMK m the togging accoun
F. Update the CMK key policy to provide access to the logging account and business unit accounts Manually rotate the CMK every 365 days.
G. Use an AWS managed CMK in the togging account Update the CMK key policy to provide access to the togging account onl
H. Enable automatic rotation of the CMK.

Answer: A

NEW QUESTION 182
- (Exam Topic 2)
A large company recently experienced an unexpected increase in Amazon RDS and Amazon DynamoDB costs The company needs to increase visibility into details of AWS Billing and Cost Management There are various accounts associated with AWS Organizations, including many development and production accounts. There is no consistent tagging strategy across the organization, but there are guidelines in place that require all infrastructure to be deployed using AWS Cloud Formation with consistent tagging Management requires cost center numbers and project ID numbers for all existing and future DynamoDB tables and RDS instances
Which strategy should the solutions architect provide to meet these requirements?

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

Answer: B

NEW QUESTION 186
- (Exam Topic 2)
A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances. Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances.
To meet regulatory and business requirements, the company must make the following changes for data backups:
• Backups must be retained based on custom daily, weekly, and monthly requirements.
• Backups must be replicated to at least one other AWS Region immediately after capture.
• The backup solution must provide a single source of backup status across the AWS environment.
• The backup solution must send immediate notifications upon failure of any resource backup.
Which combination of steps will meet these requirements with the LEAST amount of operational overhead? (Select THREE.)

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

Answer: BDE

NEW QUESTION 189
- (Exam Topic 2)
A company owns a chain of travel agencies and is running an application in the AWS Cloud. Company employees use the application to search (or Information about travel destinations. Destination content is updated four times each year.
Two fixed Amazon EC2 instances serve the application. The company uses an Amazon Route 53 public hosted zone with a multivalue record of travel.example.com that returns the Elastic IP addresses for the EC2 instances. The application uses Amazon DynamoDB as its primary data store. The company uses a self-hosted Redis instance as a caching solution.
During content updates, the load on the EC2 instances and the caching solution increases drastically. This increased load has led to downtime on several occasions. A solutions architect must update the application so that the application is highly available and can handle the load that is generated by the content updates.
Which solution will meet these requirements?

A. Set up DynamoDB Accelerator (DAX} as in-memory cach
B. Update the application to use DA
C. Create an Auto Scaling group for the EC2 instance

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

Answer: B

NEW QUESTION 193
- (Exam Topic 2)
A company is running an application in the AWS Cloud. The company's security team must approve the creation of all new IAM users. When a new 1AM user is created, all access for the user must be removed automatically. The security team must then receive a notification to approve the user. The company has a multi-Region AWS CloudTrail trail In the AWS account.
Which combination of steps will meet these requirements? (Select THREE.)

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rul
B. Define a pattern with the detail-type value set to AWS API Call via CloudTrail and an eventName of CreateUser.
C. Configure CloudTrail to send a notification for the CreateUser event to an Amazon Simple Notification Service (Amazon SNS) topic.
D. Invoke a container that runs in Amazon Elastic Container Service (Amazon ECS) with AWS Fargatetechnology to remove access
E. Invoke an AWS Step Functions state machine to remove access.
F. Use Amazon Simple Notification Service (Amazon SNS) to notify the security team.
G. Use Amazon Pinpoint to notify the security team.

Answer: ABE

NEW QUESTION 197
- (Exam Topic 2)
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-line 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 complexly.
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 for remote access
E. Use AWS KMS to store database credentials
F. Use AWS Systems Manager Session Manager for remote access

Answer: ACF

NEW QUESTION 201
- (Exam Topic 2)
A company is running its solution on AWS in a manually created VPC. The company is using AWS Cloud Formation to provision other parts of the infrastructure. According to a new requirement, the company must manage all infrastructure in an automatic way.
What should the company do to meet this new requirement with the LEAST effort?

A. Create a new AWS Cloud Development Kit (AWS CDK) stack that stnctly provisions the existing VPC resources and configuratio
B. Use AWS CDK to import the VPC into the stack and to manage the VPC.
C. Create a CloudFormation stack set that creates the VP
D. Use the stack set to import the VPC into the stack.
E. Create a new CloudFormation template that strictly provisions the existing VPC resources and configuratio
F. From the CloudFormation console, create a new stack by importing the existing resources.
G. Create a new CloudFormation template that creates the VP
H. Use the AWS Serverless Application Model {AWS SAM) CLI to import the VPC.

Answer: D

NEW QUESTION 204
- (Exam Topic 2)
A company is planning to migrate an application from on premises to AWS. The application currently uses an Oracle database and the company can tolerate a brief downtime of 1 hour when performing the switch to the new infrastructure As part of the migration. the database engine will be changed to MySQL. A solutions architect needs to determine which AWS services can be used to perform the migration while minimizing the amount of work and time required.

Which of the following will meet the requirements?

A. Use AWS SCT to generate the schema scripts and apply them on the target prior to migration Use AWS DMS to analyse the current schema and provide a recommendation for the optimal database engine Then, use AWS DMS to migrate to the recommended engine Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually
B. Use AWS SCT to generate the schema scripts and apply them on the target prior to migratio
C. Use AWS DMS to begin moving data from the on-premises database to AW
D. After the initial copy continue to use AWS DMS to keep the databases m sync until cutting over to the new database Use AWS SCT to identify what embedded SOL code in the application can be converted and what has to be done manually.
E. Use AWS DMS lo help identify the best target deployment between installing the database engine on Amazon EC2 directly or moving to Amazon RD
F. Then, use AWS DMS to migrate to the platfor
G. Use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually.
H. Use AWS DMS to begin moving data from the on-premises database to AWS After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new database use AWS Application Discovery Service to identify what embedded SQL code m the application can be convened and what has to be done manually

Answer: B

NEW QUESTION 207
- (Exam Topic 2)
A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts Each VPC consists of public subnets and private subnets that span across multiple Availability Zones NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets.
A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account
Which set of additional steps should the solutions architect take to meet these requirements?

A. Create peering connections between the egress VPC and the spoke VPCs Configure the required routing to allow access to the internet
B. Create a transit gateway and share it with the existing AWS accounts Attach existing VPCs to the transit gateway Configure the required routing to allow access to the internet
C. Create a transit gateway in every account Attach the NAT gateway to the transit gateways Configure the required routing to allow access to the internet
D. Create an AWS PrivateLink connection between the egress VPC and the spoke VPCs Configure the required routing to allow access to the internet

Answer: B

NEW QUESTION 209
- (Exam Topic 2)
A company has an organization that has many AWS accounts in AWS Organizations A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.
The company has a common set of IP CIDR ranges in an allow list in each AWS account lo allow access to and from the company's on-premises network Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.
The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts Which solution meets these requirements with the LEAST amount of operational overhead.

A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS account Deploy an AWS Lambda function in each AWS account Configure the Lambda function to run every time an SNS topic receives a message Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account Instruct the security team to distribute changes by publishing messages to its SNS topic
B. Create new customer-managed prefix lists in each AWS account within the organization Populate theprefix lists in each account with all internal CIDR ranges Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups Instruct the security team to share updates with each AWS account owner.
C. Create a new customer-managed prefix list in the security team's AWS account Populate thecustomer-managed prefix list with all internal CIDR range
D. Share the customer-managed prefix list.... organization by using AWS Resource Access Manager Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups

Answer: A

NEW QUESTION 212
......

Relate Links

100% Pass Your SAP-C02 Exam with Exambible Prep Materials

https://www.exambible.com/SAP-C02-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/