Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
- (Exam Topic 1)
A company standardized its method of deploying applications to AWS using AWS CodePipeline and AWS Cloud Formation. The applications are in Typescript and Python. The company has recently acquired another business that deploys applications to AWS using Python scripts.
Developers from the newly acquired company are hesitant to move their applications under CloudFormation because it would require than they learn a new domain-specific language and eliminate their access to language features, such as looping.
How can the acquired applications quickly be brought up to deployment standards while addressing the developers' concerns?

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

Answer: D

NEW QUESTION 2
- (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 3
- (Exam Topic 1)
A solution architect needs to deploy an application on a fleet of Amazon EC2 instances. The EC2 instances run in private subnets in An Auto Scaling group. The application is expected to generate logs at a rate of 100 MB each second on each of the EC2 instances.
The logs must be stored in an Amazon S3 bucket so that an Amazon EMR cluster can consume them for further processing The logs must be quickly accessible for the first 90 days and should be retrievable within 48 hours thereafter.
What is the MOST cost-effective solution that meets these requirements?

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

Answer: C

NEW QUESTION 4
- (Exam Topic 1)
A company built an ecommerce website on AWS using a three-tier web architecture. The application is
Java-based and composed of an Amazon CloudFront distribution, an Apache web server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database.
Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts. The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics. Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis.
Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events? (Select THREE.)

A. Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs.
B. Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java.
C. Configure the Aurora MySQL DB cluster to stream slow query and error logs to Amazon Kinesis.

D. Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs.
E. Enable and configure AWS CloudTrail to collect and analyze application activity from Amazon EC2 and Aurora.
F. Enable Aurora MySQL DB cluster performance benchmarking and publish the stream to AWS X-Ray.

Answer: ABD

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_LogAccess.Concepts.MySQL.html# https://aws.amazon.com/blogs/mt/simplifying-apache-server-logs-with-amazon-cloudwatch-logs-insights/ https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-dotnet-messagehandler.html https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-java-sqlclients.html

NEW QUESTION 5
- (Exam Topic 1)
A company that is developing a mobile game is making game assets available in two AWS Regions. Game assets ate served from a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in each Region. The company requires game assets to be (etched from the closest Region. If game assets become unavailable in the closest Region, they should be fetched from the other Region.
What should a solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distributio
B. Create an origin group with one origin for each AL
C. Set one of the origins as primary.
D. Create an Amazon Route 53 health check for each AL
E. Create a Route 53 failover routing record pointing to the two ALB
F. Set the Evaluate Target Health value to Yes.
G. Create two Amazon CloudFront distributions, each with one ALB as the origi
H. Create an Amazon Route 53 failover routing record pointing to the two CloudFront distribution
I. Set the Evaluate Target Health value to Yes.
J. Create an Amazon Route 53 health check for each AL
K. Create a Route 53 latency alias record pointing to the two ALB
L. Set the Evaluate Target Health value to Yes.

Answer: D

Explanation:
Failover routing policy – Use when you want to configure active-passive failover. Latency routing policy – Use when you have resources in multiple AWS Regions and you want to route traffic to the region that provides the best latency. https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

NEW QUESTION 6
- (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 7
- (Exam Topic 1)
A company is building an image service on the web that will allow users to upload and search random photos. At peak usage, up to 10.000 users worldwide will upload their images. The service will then overlay text on the uploaded images, which will then be published on the company website.
Which design should a solutions architect implement?

A. Store the uploaded images in Amazon Elastic File System (Amazon EFS). Send application log information about each image to Amazon CloudWatch Log
B. Create a fleet of Amazon EC2 instances that use CloudWatch Logs to determine which images need to be processe
C. Place processed images in anolher directory in Amazon EF
D. Enable Amazon CloudFront and configure the origin to be the one of the EC2 instances in the fleet.
E. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to Amazon Simple Notification Service (Amazon SNS). Create a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) to pull messages from Amazon SNS to process the images

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

Answer: C

NEW QUESTION 8
- (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 9
- (Exam Topic 1)
A large company with hundreds of AWS accounts has a newly established centralized internal process for purchasing new or modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement or execution. Previously, business units would directly purchase or modify Reserved Instances in their own respective AWS accounts autonomously.
Which combination of steps should be taken to proactively enforce the new process in the MOST secure way possible? (Select TWO.)

A. Ensure all AWS accounts are part of an AWS Organizations structure operating in all features mode.
B. Use AWS Contig lo report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedlnstancesOffering and ec2:ModifyReservedlnstances actions.
C. In each AWS account, create an IAM policy with a DENY rule to the ec2:PurchaseReservedlnstancesOffering and ec2:ModifyReservedInstances actions.
D. Create an SCP that contains a deny rule to the ec2:PurchaseReservedlnstancesOffering and ec2: Modify Reserved Instances action
E. Attach the SCP to each organizational unit (OU) of the AWS Organizations structure.
F. Ensure that all AWS accounts are part of an AWS Organizations structure operating in consolidated billing features mode.

Answer: AD

Explanation:
https://docs.aws.amazon.com/organizations/latest/APIReference/API_EnableAllFeatures.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp-strategies.html

NEW QUESTION 10
- (Exam Topic 1)
A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency network connections between all of the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 10
- (Exam Topic 1)
A 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 the Amazon 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 11
- (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 12
- (Exam Topic 1)
A company has multiple AWS accounts as part of an organization created with AWS Organizations. Each account has a VPC in the us-east-2 Region and is used for either production or development workloads. Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other.
To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another.
Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

A. Modify the security groups assigned to development EC2 instances to block traffic from production EC2 instance
B. Modify the security groups assigned to production EC2 instances to block traffic from development EC2 instances.
C. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attache
D. Using the Network Manager feature of AWS Transit Gateway, create policies that restrict traffic between VPCs based on the value of this tag.
E. Create separate route tables for production and development traffi
F. Delete each account's association and route propagation to the default AWS Transit Gateway route tabl
G. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment.
H. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attache
I. Modify the AWS Transit Gateway routing table to route production tagged attachments to one another and development tagged attachments to one another.

Answer: C

Explanation:
https://docs.aws.amazon.com/vpc/latest/tgw/vpc-tgw.pdf

NEW QUESTION 17
- (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 21
- (Exam Topic 1)

A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.
The company uses Amazon Elastic Container Registry (Amazon ECRJ to store its container images When a new image version is uploaded, the new image version receives a unique tag
The company needs a solution that inspects new image versions for common vulnerabilities and exposures The solution must automatically delete new image tags that have Critical or High severity findings The solution also must notify the development team when such a deletion occurs
Which solution meets these requirements?

A. Configure scan on push on the repositor
B. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical or High severity findings Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS)
C. Configure scan on push on the repository Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queue Invoke an AWS Lambda function when a new message is added to the SOS queue Use the Lambda function to delete the image tag for images that have Critical or High seventy finding
D. Notify the development team by using Amazon Simple Email Service (Amazon SES).
E. Schedule an AWS Lambda function to start a manual image scan every hour Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke another Lambda function when a scan is complet
F. Use the second Lambda function to delete the image tag for images that have Cnocal or High severity finding
G. Notify the development team by using Amazon Simple Notification Service (Amazon SNS)
H. Configure periodic image scan on the repository Configure scan results to be added to an Amazon Simple Queue Service (Amazon SQS) queue Invoke an AWS Step Functions state machine when a new message is added to the SQS queue Use the Step Functions state machine to delete the image tag for images that have Critical or High severity finding
I. Notify the development team by using Amazon Simple Email Service (Amazon SES).

Answer: C

NEW QUESTION 22
- (Exam Topic 1)
A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.
Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instanc
B. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling grou
C. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.
D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling grou
E. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.
F. Change the log delivery rate to every 5 minute
G. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user dat
H. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance terminatio
I. Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.
J. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service(Amazon SNS) topi
K. From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Answer: B

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/adding-lifecycle-hooks.html
- Refer to Default Result section - If the instance is terminating, both abandon and continue allow the instance to terminate. However, abandon stops any remaining actions, such as other lifecycle hooks, and continue allows any other lifecycle hooks to complete.
https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-i https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function
https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function/blob/master/cloudformation/template.yam

NEW QUESTION 24
- (Exam Topic 1)
A solutions architect is evaluating the reliability of a recently migrated application running on AWS. The front end is hosted on Amazon S3 and accelerated by Amazon CloudFront. The application layer is running in a stateless Docker container on an Amazon EC2 On-Demand Instance with an Elastic IP address. The storage layer is a MongoDB database running on an EC2 Reserved Instance in the same Availability Zone as the application layer.
Which combination of steps should the solutions architect take to eliminate single points of failure with minimal application code changes? (Select TWO.)

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

Answer: BD

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

NEW QUESTION 26
- (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 in the 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 28
- (Exam Topic 1)
A company has 50 AWS accounts that are members of an organization in AWS Organizations Each account contains multiple VPCs The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.
Which combination of steps will meet these requirements? (Select TWO)

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

Answer: AC

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

NEW QUESTION 32
- (Exam Topic 1)
A financial services company receives a regular data feed from its credit card servicing partner Approximately 5.1 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.
Which solutions will meet these requirements?

A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queu
B. Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
C. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queu
D. Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains message
E. Have the application process each record, and transform the record into JSON forma
F. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
G. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to matc
H. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirement
I. Define the output format as JSO
J. Once complete, have the ETL job send the results to another S3 bucket for internal processing.
K. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to matc
L. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirement
M. Define the output format as JSO
N. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Answer: C

Explanation:
You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object create events occur. The Lambda function will then trigger the Glue ETL job to transform the records masking the sensitive data and modifying the output format to JSON. This solution meets all requirements.
Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON.
Once complete, have the ETL job send the results to another S3 bucket for internal processing. https://docs.aws.amazon.com/glue/latest/dg/trigger-job.html https://d1.awsstatic.com/Products/product-name/diagrams/product-page-diagram_Glue_Event-driven-ETL-Pipel

NEW QUESTION 33
- (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 37
- (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 accountwhen the logs are being copied.

Answer: B

NEW QUESTION 41
- (Exam Topic 1)
A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.
The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.
Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?

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

Answer: A

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

NEW QUESTION 46

- (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 51
- (Exam Topic 1)
A company uses AWS Transit Gateway for a hub-and-spoke model to manage network traffic between many VPCs. The company is developing a new service that must be able to send data at 100 Gbps. The company needs a faster connection to other VPCs in the same AWS Region.
Which solution will meet these requirements?

A. Establish VPC peering between the necessary VPC
B. Ensure that all route tables are updated as required.
C. Attach an additional transit gateway to the VPC
D. Update the route tables accordingly.
E. Create AWS Site-to-Site VPN connections that use equal-cost multi-path (ECMP) routing between the necessary VPCs.
F. Create an additional attachment from the necessary VPCs to the existing transit gateway.

Answer: D

NEW QUESTION 56
- (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 59
- (Exam Topic 1)
A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand.
Which solutions meet these requirements? (Choose two.)

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

Answer: CD

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

NEW QUESTION 61
- (Exam Topic 1)
A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services. Database credentials are hard-coded on each instance. SSH keys for command-tine remote access are stored in a secured Amazon S3 bucket. The company has asked its solutions architect to improve the security posture of the architecture without adding operational complexity.
Which combination of steps should the solutions architect take to accomplish this? (Select THREE.)

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

Answer: ACF

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

NEW QUESTION 66
- (Exam Topic 1)
A solutions architect is designing a publicly accessible web application that is on an Amazon CloudFront distribution with an Amazon S3 website endpoint as the origin. When the solution is deployed, the website returns an Error 403: Access Denied message.
Which steps should the solutions architect take to correct the issue? (Select TWO.)

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

Answer: AB

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

NEW QUESTION 70
- (Exam Topic 1)
A company hosts a web application that tuns on a group of Amazon EC2 instances that ate behind an Application Load Balancer (ALB) in a VPC. The company wants to analyze the network payloads lo reverse-engineer a sophisticated attack of the application.
Which approach should the company take to achieve this goal?

A. Enable VPC Flow Log
B. Store the flow logs in an Amazon S3 bucket for analysis.
C. Enable Traffic Mirroring on the network interface of the EC2 instance
D. Send the mirrored traffic lo a target for storage and analysis.
E. Create an AWS WAF web AC
F. and associate it with the AL
G. Configure AWS WAF logging.
H. Enable logging for the AL
I. Store the logs in an Amazon S3 bucket for analysis.

Answer: A

NEW QUESTION 71
- (Exam Topic 1)
A company uses an on-premises data analytics platform. The system is highly available in a fully redundant configuration across 12 servers in the company's data center.
The system runs scheduled jobs, both hourly and daily, in addition to one-time requests from users. Scheduled jobs can take between 20 minutes and 2 hours to finish running and have tight SLAs. The scheduled jobs account for 65% of the system usage. User jobs typically finish running in less than 5 minutes and have no SLA. The user jobs account for 35% of system usage. During system failures, scheduled jobs must continue to meet SLAs. However, user jobs can be delayed.
A solutions architect needs to move the system to Amazon EC2 instances and adopt a consumption-based model to reduce costs with no long-term commitments. The solution must maintain high availability and must not affect the SLAs.
Which solution will meet these requirements MOST cost-effectively?

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

I. Run two instances in each Availability Zone as Spot Instances.
J. Split the 12 instances across three Availability Zones in the chosen AWS Regio
K. Run three instances in each Availability Zone as On-Demand Instances with Capacity Reservation
L. Run one instance in each Availability Zone as a Spot Instance.

Answer: D

NEW QUESTION 73
- (Exam Topic 1)
A solutions architect is designing an application to accept timesheet entries from employees on their mobile devices. Timesheets will be submitted weekly, with most of the submissions occurring on Friday. The data must be stored in a format that allows payroll administrators to run monthly reports. The infrastructure must be highly available and scale to match the rate of incoming data and reporting requests.
Which combination of steps meets these requirements while minimizing operational overhead? (Select TWO.)

A. Deploy the application to Amazon EC2 On-Demand Instances With load balancing across multiple Availability Zone
B. Use scheduled Amazon EC2 Auto Scaling to add capacity before the high volume of submissions on Fridays.
C. Deploy the application in a container using Amazon Elastic Container Service (Amazon ECS) with load balancing across multiple Availability Zone
D. Use scheduled Service Auto Scaling to add capacity before the high volume of submissions on Fridays.
E. Deploy the application front end to an Amazon S3 bucket served by Amazon CloudFron
F. Deploy the application backend using Amazon API Gateway with an AWS Lambda proxy integration.
G. Store the timesheet submission data in Amazon Redshif
H. Use Amazon OuickSight to generate the reports using Amazon Redshift as the data source.
I. Store the timesheet submission data in Amazon S3. Use Amazon Athena and Amazon OuickSight to generate the reports using Amazon S3 as the data source.

Answer: AE

NEW QUESTION 78
- (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 80
- (Exam Topic 1)
A company wants to retire its Oracle Solaris NFS storage arrays. The company requires rapid data migration over its internet network connection to a combination of destinations for Amazon S3. Amazon Elastic File System (Amazon EFS), and Amazon FSx lor Windows File Server. The company also requires a full initial copy, as well as incremental transfers of changes until the retirement of the storage arrays. All data must be encrypted and checked for integrity.
What should a solutions architect recommend to meet these requirements?

A. Configure CloudEndur
B. Create a project and deploy the CloudEndure agent and token to the storage arra
C. Run the migration plan to start the transfer.
D. Configure AWS DataSyn
E. Configure the DataSync agent and deploy it to the local networ
F. Create a transfer task and start the transfer.
G. Configure the aws S3 sync comman
H. Configure the AWS client on the client side with credential
I. Run the sync command to start the transfer.
J. Configure AWS Transfer (or FT
K. Configure the FTP client with credential
L. Script the client to connect and sync to start the transfer.

Answer: B

NEW QUESTION 82
- (Exam Topic 1)
A company is planning on hosting its ecommerce platform on AWS using a multi-tier web application designed for a NoSQL database. The company plans to use the us-west-2 Region as its primary Region. The company want to ensure that copies of the application and data are available in a second Region, us-west-1, for disaster recovery. The company wants to keep the time to fail over as low as possible. Failing back to the primary Region should be possible without administrative interaction after the primary service is restored.
Which design should the solutions architect use?

A. Use AWS Cloud Formation StackSets lo create the stacks in both Regions with Auto Scaling groups for the web and application tier
B. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replicatio
C. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outag
D. Use Amazon DynamoDB global tables for the database tier.

E. Use AWS Cloud Formation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tier
F. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replicatio
G. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outag
H. Deploy an Amazon Aurora global database for the database tier.
I. Use AWS Service Catalog to deploy the web and application servers in both Region
J. Asynchronously replicate static content between the two Regions using Amazon S3 cross-Region replicatio
K. Use Amazon Route 53 health checks to identify a primary Region failure and update the public DNS entry listing to the secondary Region in the event of an outag
L. Use Amazon RDS for MySQL withcross-Region replication for the database tier.
M. Use AWS CloudFormation StackSets to create the stacks in both Regions using Auto Scaling groups for the web and application tier
N. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replicatio
O. Use Amazon CloudFront with static files in Amazon S3, and multi-Region origins for the front-end web tie
P. Use Amazon DynamoD8 tables in each Region with scheduled backups to Amazon S3.

Answer: A

NEW QUESTION 86
- (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 91
- (Exam Topic 1)
The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.
Which combination of actions will meet these requirements? (Select THREE.)

A. Activate the user-defined cost allocation tags that represent the application and the team.
B. Activate the AWS generated cost allocation tags that represent the application and the team.
C. Create a cost category for each application in Billing and Cost Management.
D. Activate IAM access to Billing and Cost Management.
E. Create a cost budget.
F. Enable Cost Explorer.

Answer: ACF

Explanation:
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html https://aws.amazon.com/premiumsupport/knowledge-center/cost-explorer-analyze-spending-and-usage/ https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html
https://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html

NEW QUESTION 93
- (Exam Topic 1)
A company is hosting a single-page web application in the AWS Cloud. The company is using Amazon CloudFront to reach its goal audience. The CloudFront distribution has an Amazon S3 bucket that is configured as its origin. The static files for the web application are stored in this S3 bucket.
The company has used a simple routing policy to configure an Amazon Route 53 A record The record points to the CloudFront distribution The company wants to use a canary deployment release strategy for new versions of the application.
What should a solutions architect recommend to meet these requirements?

A. Create a second CloudFront distribution for the new version of the applicatio
B. Update the Route 53 record to use a weighted routing policy.
C. Create a Lambda@Edge functio
D. Configure the function to implement a weighting algorithm and rewrite the URL to direct users to a new version of the application.
E. Create a second S3 bucket and a second CloudFront origin for the new S3 bucket Create a CloudFront origin group that contains both origins Configure origin weighting for the origin group.
F. Create two Lambda@Edge function
G. Use each function to serve one of the application versions Set up a CloudFront weighted Lambda@Edge invocation policy

Answer: A

NEW QUESTION 96

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

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

Answer: B

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

NEW QUESTION 101
- (Exam Topic 1)
A company has a website that enables users to upload videos. Company policy states the uploaded videos must be analyzed for restricted content. An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SOS queue with the video's location. A backend application pulls this location from Amazon SOS and analyzes the video.
The video analysis is compute-intensive and occurs sporadically during the day The website scales with demand. The video analysis application runs on a fixed number of instances. Peak demand occurs during the holidays, so the company must add instances to the application dunng this time. All instances used are currently on-demand Amazon EC2 T2 instances. The company wants to reduce the cost of the current solution.
Which of the following solutions is MOST cost-effective?

A. Keep the website on T2 instance
B. Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to cover peak deman
C. Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application
D. Keep the website on T2 instance
E. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak deman
F. Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances.
G. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instance
H. Determine the minimum number of website instances required during off-peak times and use On-Demand Instances to cover them while using Spot capacity to cover peak demand Use Spot Fleet for the video anarysis application comprised of C4 and Amazon EC2 C5 instances.
I. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instance
J. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand Use Spot Fleet for the video analysis application comprised of R4 and Amazon EC2 R5 instances

Answer: B

NEW QUESTION 104
- (Exam Topic 1)
A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files ate uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.examWe.com through the use of Amazon Route 53.
What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

A. Move the EC2 instance into an Auto Scaling grou
B. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record sftp.example.com in Route 53 to point to the ALB.
C. Migrate the SFTP server to AWS Transfer for SFT
D. Update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname.
E. Migrate the SFTP server to a file gateway in AWS Storage Gatewa
F. Update the DNS record sflp.example.com in Route 53 to point to the file gateway endpoint.
G. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record sftp.example.com in Route 53 to point to the NLB.

Answer: B

Explanation:
https://aws.amazon.com/aws-transfer-family/faqs/ https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html
https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/?nc1=h_

NEW QUESTION 107
- (Exam Topic 1)
A company is planning to set up a REST API application on AWS. The application team wants to set up a new identity store on AWS The IT team does not want to maintain any infrastructure or servers for this deployment.
What is the MOST operationally efficient solution that meets these requirements?

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

C. Deploy the application in AWS AppSync, and configure AWS Lambda resolvers Set up an Amazon Cognito user pool, and configure AWS AppSync to use the user pool for authorization
D. Deploy the application as AWS Lambda function
E. Set up Amazon API Gateway REST API endpoints for the application Set up an Amazon Cognito user pool, and configure an Amazon Cognito authorizer
F. Deploy the application in Amazon Elastic Kubemetes Service (Amazon EKS) cluster
G. Set up an Application Load Balancer for the EKS pods Set up an Amazon Cognito user pool and service pod for authentication.

Answer: C

NEW QUESTION 108
- (Exam Topic 1)
A company is running a tone-of-business (LOB) application on AWS to support its users The application runs in one VPC. with a backup copy in a second VPC in a different AWS Region for disaster recovery The company has a single AWS Direct Connect connection between its on-premises network and AWS The connection terminates at a Direct Connect gateway
All access to the application must originate from the company's on-premises network, and traffic must be
encrypted in transit through the use of Psec. The company is routing traffic through a VPN tunnel over the Direct Connect connection to provide the required encryption.
A business continuity audit determines that the Direct Connect connection represents a potential single point of failure for access to the application The company needs to remediate this issue as quickly as possible.
Which approach will meet these requirements?

A. Order a second Direct Connect connection to a different Direct Connect locatio
B. Terminate the second Direct Connect connection at the same Direct Connect gateway.
C. Configure an AWS Site-to-Site VPN connection over the internet Terminate the VPN connection at a virtual private gateway in the secondary Region
D. Create a transit gateway Attach the VPCs to the transit gateway, and connect the transit gateway to the Direct Connect gateway Configure an AWS Site-to-Site VPN connection, and terminate it at the transit gateway
E. Create a transit gatewa
F. Attach the VPCs to the transit gateway, and connect the transit gateway to the Direct Connect gatewa
G. Order a second Direct Connect connection, and terminate it at the transit gateway.

Answer: C

Explanation:
Create a transit gateway. Attach the VPCs to the transit gateway, and connect the transit gateway to the Direct Connect gateway. Configure an AWS Site-to- Site VPN connection, and terminate it at the transit gateway
https://aws.amazon.com/premiumsupport/knowledge-center/dx-configure-dx-and-vpn-failover-tgw/
All access to the application must originate from the company’s on-premises network and traffic must be encrypted in transit through the use of IPsec. = need to use VPN.

NEW QUESTION 109
- (Exam Topic 1)
A solutions architect works for a government agency that has strict disaster recovery requirements All Amazon Elastic Block Store (Amazon EBS) snapshots are required to be saved in at least two additional AWS Regions. The agency also is required to maintain the lowest possible operational overhead.
Which solution meets these requirements?

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

Answer: B

NEW QUESTION 112
- (Exam Topic 1)
A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application’s data. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record.
The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an
active-active strategy.
What should a solutions architect recommend to meet these requirements?

A. Reconfigure the application’s Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
C. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Regio
D. Configure the CloudWatch alarm to invoke the Lambda function.
E. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
F. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealth
G. Update the application’s Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.
H. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Regio
I. Reconfigure the application’s Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
J. Remove the read replic
K. Replace the read replica with a standalone RDS DB instanc
L. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.
M. Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted target
N. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
O. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Regio

P. Configure the CloudWatch alarm to invoke the Lambda function.

Answer: B

Explanation:
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

NEW QUESTION 117
- (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 metadata search 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 118
- (Exam Topic 1)
A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The maximum size ol the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:
VPC CIDR: 10.0.0.0/23
AZ1 subnet CIDR: 10.0.0.0/24 AZ2 subnet CIDR: 10.0.1.0/24
Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 120
- (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 extremety 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 125
- (Exam Topic 1)
A company Is serving files to its customers through an SFTP server that Is accessible over the internet The SFTP server Is running on a single Amazon EC2 instance with an Elastic IP address attached Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication The EC2 instance also has an attached security group that allows access from all customer IP addresses.
A solutions architect must implement a solution to improve availability minimize the complexity ot infrastructure management and minimize the disruption to customers who access files. The solution must not change the way customers connect.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 130
- (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: D

NEW QUESTION 134
- (Exam Topic 1)
A company plans to migrate to AWS. A solutions architect uses AWS Application Discovery Service over the fleet and discovers that there is an Oracle data warehouse and several PostgreSQL databases. Which combination of migration patterns will reduce licensing costs and operational overhead? (Select TWO.)

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

Answer: BD

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

NEW QUESTION 138
- (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 142
- (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 143
- (Exam Topic 1)
A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization

in AWS Organizations. The company requires the cost lor cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.
Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)

A. Create an AWS Config rule in each account to find resources with missing tags.
B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
C. Use Amazon Inspector in the organization to find resources with missing tags.
D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.

Answer: CDE

NEW QUESTION 146
- (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 149
- (Exam Topic 1)
A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.
Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

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

Answer: C

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

NEW QUESTION 151
- (Exam Topic 1)
An online retail company hosts its stateful web-based application and MySQL database in an on-premises data center on a single server. The company wants to increase its customer base by conducting more marketing campaigns and promotions. In preparation, the company wants to migrate its application and database to AWS to increase the reliability of its architecture.
Which solution should provide the HIGHEST level of reliability?

A. Migrate the database to an Amazon RDS MySQL Multi-AZ DB instanc
B. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
C. Store sessions in Amazon Neptune.
D. Migrate the database to Amazon Aurora MySQ
E. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
F. Store sessions in an Amazon ElastiCache for Redis replication group.
G. Migrate the database to Amazon DocumentDB (with MongoDB compatibility). Deploy the application in an Auto Scaling group on Amazon EC2 instances behind a Network Load Balance
H. Store sessions in Amazon Kinesis Data Firehose.
I. Migrate the database to an Amazon RDS MariaDB Multi-AZ DB instanc
J. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
K. Store sessions in Amazon ElastiCache for Memcached.

Answer: B

NEW QUESTION 153
- (Exam Topic 1)
A company has a policy that all Amazon EC2 instances that are running a database must exist within the same subnets in a shared VPC Administrators must follow security compliance requirements and are not allowed to directly log in to the shared account All company accounts are members of the same organization in AWS Organizations. The number of accounts will rapidly increase as the company grows.
A solutions architect uses AWS Resource Access Manager to create a resource share in the shared account What is the MOST operationally efficient configuration to meet these requirements?

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

Answer: C

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

NEW QUESTION 158
- (Exam Topic 1)
To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network. The security team mandates that no traffic accessing this data should traverse the public internet.
How should the solutions architect design a highly available solution that meets the requirements and is cost-effective?

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

Answer: D

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

NEW QUESTION 159
- (Exam Topic 1)
A financial services company logs personally identifiable information 10 its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The security team has mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.
Which steps should the solutions architect take to meet these requirements?

A. Create an AWS CloudHSM cluste
B. Create a new CMK in AWS KMS using AWS_CloudHSM as the source (or the key material and an origin of AWS_CLOUDHS
C. Enable automatic key rotation on the CMK with a duration of 1 yea
D. Configure a bucket policy on the togging bucket thai disallows uploads of unencrypted data and requires that the encryption source be AWS KMS.
E. Provision an AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VPC
F. Configure an AWS bucket policy on the logging bucket that requires all objects to be encrypte
G. Configure the logging application to query theon-premises HSMs from the AWS environment for the encryption key material, and create a unique CMK for each logging event.
H. Create a CMK in AWS KMS with no key material and an origin of EXTERNA
I. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AW
J. Configure a bucket policy on the logging bucket that disallows uploads ofnon-encrypted data and requires that the encryption source be AWS KMS.
K. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS_KMS.Disable this CM
L. and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AW
M. Re-enable the CM
N. Enable automatic key rotation on the CMK with a duration of 1 yea
O. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Answer: C

Explanation:
https://aws.amazon.com/blogs/security/how-to-byok-bring-your-own-key-to-aws-kms-for-less-than-15-00-a-yea
https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-create-cmk.html

NEW QUESTION 161
- (Exam Topic 1)
A company wants to migrate its corporate data center from on premises to the AWS Cloud. The data center includes physical servers and VMs that use VMware

and Hyper-V. An administrator needs to select the correct services to collect data (or the initial migration discovery process. The data format should be supported by AWS Migration Hub. The company also needs the ability to generate reports from the data.
Which solution meets these requirements?

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

Answer: C

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

NEW QUESTION 162
- (Exam Topic 1)
A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.
The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS.
Which data migration strategy should the company use?

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

Answer: B

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

NEW QUESTION 163
- (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 164
- (Exam Topic 1)
An education company is running a web application used by college students around the world. The application runs in an Amazon Elastic Container Service {Amazon ECS) cluster in an Auto Scaling group behind an Application Load Balancer (ALB). A system administrator detects a weekly spike in the number of failed login attempts, which overwhelm the application's authentication service. All the failed login attempts originate from about 500 different IP addresses that change each week, A solutions architect must prevent the failed login attempts from overwhelming the authentication service.
Which solution meets these requirements with the MOST operational efficiency?

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

Answer: B

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

NEW QUESTION 167
- (Exam Topic 1)
A web application is hosted in a dedicated VPC that is connected to a company's on-premises data center over a Site-to-Site VPN connection. The application is accessible from the company network only. This is a temporary non-production application that is used during business hours. The workload is generally low with occasional surges.
The application has an Amazon Aurora MySQL provisioned database cluster on the backend. The VPC has an internet gateway and a NAT gateways attached. The web servers are in private subnets in an Auto Scaling group behind an Elastic Load Balancer. The web servers also upload data to an Amazon S3 bucket through the internet.
A solutions architect needs to reduce operational costs and simplify the architecture. Which strategy should the solutions architect use?

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

Answer: B

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

NEW QUESTION 171
- (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 175
- (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 the correct 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 177
- (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 179
- (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 183
- (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 187
- (Exam Topic 1)
A company is building a hybrid solution between its existing on-premises systems and a new backend in AWS. The company has a management application to monitor the state of its current IT infrastructure and automate responses to issues. The company wants to incorporate the status of its consumed AWS services into the application. The application uses an HTTPS endpoint to receive updates.
Which approach meets these requirements with the LEAST amount of operational overhead?

A. Configure AWS Systems Manager OpsCenter to ingest operational events from the on-premises systems Retire the on-premises management application and adopt OpsCenter as the hub
B. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Personal Health Dashboard Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to the HTTPS endpoint of the management application
C. Modify the on-premises management application to call the AWS Health API to poll for status events of AWS services.
D. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Service Health Dashboard Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to an HTTPS endpoint for the management application with a topic filter corresponding to the services being used

Answer: A

Explanation:
ALB & NLB both supports IPs as targets. Questions is based on TCP traffic over VPN to on-premise. TCP is layer 4 and the , load balancer should be NLB. Then next questions does NLB supports loadbalcning traffic over VPN. And answer is YEs based on below URL.
https://aws.amazon.com/about-aws/whats-new/2018/09/network-load-balancer-now-supports-aws-vpn/
Target as IPs for NLB & ALB: https://aws.amazon.com/elasticloadbalancing/faqs/?nc=sn&loc=5 https://aws.amazon.com/elasticloadbalancing/application-load-balancer/

NEW QUESTION 188
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

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

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

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

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

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

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