Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

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

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

Answer: B

NEW QUESTION 2
- (Exam Topic 1)
A company 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 accountnumbers 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 Conng 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 resource
E. Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
F. Create AWS WAF rules in the management account of the organizatio
G. 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 account
H. Assume the roles by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts
I. Use AWS Control Tower to manage AWS WAF rules across accounts in the organizatio
J. 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 OU
K. Create IAM users in member accounts Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: B

NEW QUESTION 3
- (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 4
- (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 5
- (Exam Topic 1)
An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory requirement for out-of-region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles.
Which of the following options can the solutions architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

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

Answer: C

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

NEW QUESTION 6
- (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_KM
L. Disable this CM
M. and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AW
N. Re-enable the CM
O. Enable automatic key rotation on the CMK with a duration of 1 yea
P. 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 7
- (Exam Topic 1)
A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability.
Which service will meet the requirements for storing the session information in the MOST cost-effective way?

A. Amazon ElastiCache with the Memcached engine
B. Amazon S3
C. Amazon RDS MySQL
D. Amazon ElastiCache with the Redis engine

Answer: D

Explanation:
https://aws.amazon.com/caching/session-management/
Building real-time apps across versatile use cases like gaming, geospatial service, caching, session stores, or queuing, with advanced data structures, replication, and point-in-time snapshot support. Memcached: Building a simple, scalable caching layer for your data-intensive apps. https://aws.amazon.com/elasticache/

NEW QUESTION 8
- (Exam Topic 1)
A company has many AWS accounts and uses AWS Organizations to manage all of them. A solutions architect must implement a solution that the company can use to share a common network across multiple accounts.
The company's infrastructure team has a dedicated infrastructure account that has a VPC. The infrastructure team must use this account to manage the network. Individual accounts cannot have the ability to manage their own networks. However, individual accounts must be able to create AWS resources within subnets.
Which combination of actions should the solutions architect perform to meet these requirements? (Select TWO.)

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

Answer: CE

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

NEW QUESTION 9
- (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 10
- (Exam Topic 1)
A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.
The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.
Which storage strategy is the MOST cost-effective and meets the design requirements?

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

Answer: B

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

NEW QUESTION 10
- (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 12
- (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 16
- (Exam Topic 1)
A company hosts a photography website on AWS that has global visitors. The website has experienced steady increases in traffic during the last 12 months, and users have reported a delay in displaying images. The company wants to configure Amazon CloudFront lo deliver photos to visitors with minimal latency.
Which actions will achieve this goal? (Select TWO.)

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

Answer: BD

NEW QUESTION 19
- (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 21
- (Exam Topic 1)
A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations lo manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.
Which solution is the MOST cost-effective way to meet these requirements?

A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owne
B. Add each business unit to an Amazon SNS topic for each aler
C. Use Cost Explorer in each account to create monthly reports for each business unit.
D. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owne
E. Add each business unit to an Amazon SNS topic for each aler
F. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
G. Configure AWS Budgets in each account and configure budget alerts lhat are grouped by application, environment, and owne
H. Add each business unit to an Amazon SNS topic for each aler
I. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.
J. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owne

K. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

Explanation:
Configure AWS Budgets in the organization€™s master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization€™s master account to create monthly reports for each business unit.
https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Bud

NEW QUESTION 25
- (Exam Topic 1)
A company has an Amazon VPC that is divided into a public subnet and a pnvate subnet. A web application runs in Amazon VPC. and each subnet has its own NACL The public subnet has a CIDR of 10.0.0 0/24 An Application Load Balancer is deployed to the public subnet The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet
Onty network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets
What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

A. An inbound rule for port 80 from source 0.0 0.0/0
B. An inbound rule for port 80 from source 10.0 0 0/24
C. An outbound rule for port 80 to destination 0.0.0.0/0
D. An outbound rule for port 80 to destination 10.0.0.0/24
E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

Answer: BE

Explanation:
Ephemeral ports are not covered in the syllabus so be careful that you don't confuse day to day best practise with what is required for the exam. Link to an explanation on Ephemeral ports here. https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-KUbcwo4lXefMl7janaK/netw

NEW QUESTION 27
- (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 AWS Organizations 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

NEW QUESTION 32
- (Exam Topic 1)
A company wants to host a new global website that consists of static content. A solutions architect is working on a solution that uses Amazon CloudFront with an origin access identity <OAI) to access website content that is stored in a private Amazon S3 bucket.
During testing, the solutions architect receives 404 errors from the S3 bucket. Error messages appear only for attempts to access paths that end with a forward slash. such as example.com/path/. These requests should return the existing S3 object path/index.html. Any potential solution must not prevent CloudFront from caching the content.
What should the solutions architect do to resolve this problem?

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

Answer: C

NEW QUESTION 37
- (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 40
- (Exam Topic 1)
A scientific organization requires the processing of text and picture data stored in an Amazon S3 bucket. The data is gathered from numerous radar stations during a mission's live, time-critical phase. The data is uploaded by the radar stations to the source S3 bucket. The data is preceded with the identification number of the radar station.
In a second account, the business built a destination S3 bucket. To satisfy a compliance target, data must be transferred from the source S3 bucket to the destination S3 bucket. Replication is accomplished by using an S3 replication rule that covers all items in the source S3 bucket.
A single radar station has been recognized as having the most precise data. At this radar station, data replication must be completed within 30 minutes of the radar station uploading the items to the source S3 bucket.
What actions should a solutions architect take to ensure that these criteria are met?

A. Set up an AWS DataSync agent to replicate the prefixed data from the source S3 bucket to the destination S3 bucke
B. Select to use at available bandwidth on the task, and monitor the task to ensure that it is in the TRANSFERRING statu
C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
D. In the second account, create another S3 bucket to receive data from the radar station with the most accurate data Set up a new replication rule for this new S3 bucket to separate the replication from the other radar stations Monitor the maximum replication time to the destinatio
E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold
F. Enable Amazon S3 Transfer Acceleration on the source S3 bucket, and configure the radar station with the most accurate data to use the new endpoint Monitor the S3 destination bucket's TotalRequestLatency metric Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes
G. Create a new S3 replication rule on the source S3 bucket that filters for the keys that use the prefix of the radar station with the most accurate data Enable S3 Replication Time Control (S3 RTC) Monitor the maximum replication time to the destination Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold

Answer: D

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

NEW QUESTION 41
- (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 42
- (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:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
C. In each AWS account, create an IAM policy with a DENY rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
D. Create an SCP that contains a deny rule to the ec2:PurchaseReservedInstancesOffering 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 47

- (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 50
- (Exam Topic 1)
An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a solutions architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible (or receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during markeling campaigns to process the orders with minimal delays.
Which of the following is the MOST reliable approach to meet the requirements?

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

Answer: B

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

NEW QUESTION 51
- (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 56
- (Exam Topic 1)
A company runs an e-commerce platform with front-end and e-commerce tiers. Both tiers run on LAMP stacks with the front-end instances running behind a load balancing appliance that has a virtual offering on AWS Current*/, the operations team uses SSH to log in to the instances to maintain patches and address other concerns. The platform has recently been the target of multiple attacks, including.

• A DDoS attack.
• An SOL injection attack
• Several successful dictionary attacks on SSH accounts on the web servers
The company wants to improve the security of the e-commerce platform by migrating to AWS. The company's solutions architects have decided to use the following approach;
• Code review the existing application and fix any SQL injection issues.
• Migrate the web application to AWS and leverage the latest AWS Linux AMI to address initial security patching.
• Install AWS Systems Manager to manage patching and allow the system administrators to run commands on all instances, as needed.
What additional steps will address all of the identified attack types while providing high availability and minimizing risk?

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

Answer: B

NEW QUESTION 57
- (Exam Topic 1)
A company is using AWS Organizations lo manage multiple accounts. Due to regulatory requirements, the company wants to restrict specific member accounts to certain AWS Regions, where they are permitted to deploy resources. The resources in the accounts must be tagged, enforced based on a group standard, and centrally managed with minimal configuration.
What should a solutions architect do to meet these requirements?

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

Answer: D

NEW QUESTION 62
- (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 65

- (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 SOS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SOS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
C. Tigger an AWS Lambda function on file delivery that extracts each record and wntes it to an Amazon SOS queu
D. Configure an AWS Fargate container application to
E. automatically scale to a single instance when the SOS queue contains message
F. Have the application process each record, and transform the record into JSON forma
G. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
H. 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 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 70
- (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 72
- (Exam Topic 1)
An education company is running a web application used by college students around the world. The application runs in an Amazon Elastic Container Service {Amazon ECS) cluster in an Auto Scaling group behind an Application Load Balancer (ALB). A system administrator detects a weekly spike in the number of failed login attempts, which overwhelm the application's authentication service. All the failed login attempts originate from about 500 different IP addresses that change each week, A solutions architect must prevent the failed login attempts from overwhelming the authentication service.
Which solution meets these requirements with the MOST operational efficiency?

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

Answer: B

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

NEW QUESTION 74
- (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 77
- (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 79
- (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 82
- (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 logsto 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 83
- (Exam Topic 1)
A company has an application that generates reports and stores them in an Amazon S3 bucket. When a user accesses their report, the application generates a signed URL to allow the user to download the report. The company's security team has discovered that the files are public and that anyone can download them without authentication. The company has suspended the generation of new reports until the problem is resolved.
Which set of actions will immediately remediate the security issue without impacting the application's normal workflow?

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

Answer: D

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

NEW QUESTION 86
- (Exam Topic 1)
A solutions architect is responsible (or redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.
Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.
Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.
Which approach would improve the availability and durability of (he system while decreasing the processing latency and minimizing costs?

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

Answer: B

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

NEW QUESTION 90
- (Exam Topic 1)
A company has developed a single-page web application in JavaScript. The source code is stored in a single Amazon S3 bucket in the us-east-1 Region. The company serves the web application to a global user base through Amazon CloudFront.
The company wants to experiment with two versions of the website without informing application users. Each version of the website will reside in its own S3 bucket. The company wants to determine which version is most successful in marketing a new product.
The solution must send application users that are based in Europe to the new website design. The solution must send application users that are based in the United States to the current website design. However, some exceptions exist. The company needs to be able to redirect specific users to the new website design, regardless of the users' location.
Which solution meets these requirements?

A. Configure two CloudFront distribution
B. Configure a geolocation routing policy in Amazon Route 53 to route traffic to the appropriate CloudFront endpoint based on the location of clients.
C. Configure a single CloudFront distributio
D. Create a behavior with different paths for each version of the sit
E. Configure Lambda@Edge on the default path to generate redirects and send the client to the correct version of the website.

F. Configure a single CloudFront distributio
G. Configure an alternate domain name on the distribution.Configure two behaviors to route users to the different S3 origins based on the domain name that the client uses in the HTTP request.
H. Configure a single CloudFront distribution with Lambda@Edg
I. Use Lambda@Edge to send user requests to different origins based on request attributes.

Answer: A

NEW QUESTION 94
- (Exam Topic 1)
A team collects and routes behavioral data for an entire company. The company runs a Multi-AZ VPC environment with public subnets, private subnets, and in internet gateway Each public subnet also contains a NAT gateway Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads run in private subnets.
A solutions architect must review the infrastructure The solutions architect needs to reduce costs and maintain the function of the applications. The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.
What should the solutions architect do to meet these requirements?

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

Answer: D

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

NEW QUESTION 95
- (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 97
- (Exam Topic 1)
A company needs to create and manage multiple AWS accounts for a number of departments from a central location. The security team requires read-only access to all accounts from its own AWs account. The company is using AWS Organizations and created an account tor the security team.
How should a solutions architect meet these requirements?

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

Answer: D

NEW QUESTION 102
- (Exam Topic 1)
A company is launching a new web application on Amazon EC2 instances. Development and production workloads exist in separate AWS accounts.

According to the company's security requirements, only automated configuration tools are allowed to access the production account. The company's security team wants to receive immediate notification if any manual access to the production AWS account or EC2 instances occurs
Which combination of actions should a solutions architect take in the production account to meet these
requirements? (Select THREE.)

A. Turn on AWS CloudTrail logs in the application's primary AWS Region Use Amazon Athena to queiy the logs for AwsConsoleSignln events.
B. Configure Amazon Simple Email Service (Amazon SES) to send email to the security team when an alarm is activated.
C. Deploy EC2 instances in an Auto Scaling group Configure the launch template to deploy instances without key pairs Configure Amazon CloudWatch Logs to capture system access logs Create an Amazon CloudWatch alarm that is based on the logs to detect when a user logs in to an EC2 instance
D. Configure an Amazon Simple Notification Service (Amazon SNS) topic to send a message to the security team when an alarm is activated
E. Turn on AWS CloudTrail logs for all AWS Region
F. Configure Amazon CloudWatch alarms to provide an alert when an AwsConsoleSignin event is detected.
G. Deploy EC2 instances in an Auto Scaling grou
H. Configure the launch template to delete the key pair after launc
I. Configure Amazon CloudWatch Logs for the system access logs Create an Amazon CloudWatch dashboard to show user logins over time.

Answer: CDE

NEW QUESTION 106
- (Exam Topic 1)
A solutions architect is designing a network for a new cloud deployment. Each account will need autonomy to modify route tables and make changes. Centralized and controlled egress internet connectivity is also needed. The cloud footprint is expected to grow to thousands of AWS accounts.
Which architecture will meet these requirements?

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

Answer: D

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

NEW QUESTION 111
- (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 113
- (Exam Topic 1)
A company maintains a restaurant review website. The website is a single-page application where files are stored in Amazon S3 and delivered using Amazon CloudFront. The company receives several fake postings every day that are manually removed.
The security team has identified that most of the fake posts are from bots with IP addresses that have a bad reputation within the same global region. The team needs to create a solution to help restrict the bots from accessing the website.
Which strategy should a solutions architect use?

A. Use AWS Firewall Manager to control the CloudFront distribution security setting
B. Create a geographical block rule and associate it with Firewall Manager.
C. Associate an AWS WAF web ACL with the CloudFront distributio
D. Select the managed Amazon IP reputation rule group for the web ACL with a deny action.
E. Use AWS Firewall Manager to control the CloudFront distribution security setting
F. Select the managed Amazon IP reputation rule group and associate it with Firewall Manager with a deny action.
G. Associate an AWS WAF web ACL with the CloudFront distributio
H. Create a rule group for the web ACL with a geographical match statement with a deny action.

Answer: B

Explanation:
IP reputation rule groups allow you to block requests based on their source. Choose one or more of these rule groups if you want to reduce your exposure to BOTS!!!! traffic or exploitation attempts
The Amazon IP reputation list rule group contains rules that are based on Amazon internal threat intelligence. This is useful if you would like to block IP addresses typically associated with bots or other threats. Inspects for a list of IP addresses that have been identified as bots by Amazon threat intelligence.

NEW QUESTION 115
- (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 116
- (Exam Topic 1)
A company manages an on-premises JavaScript front-end web application. The application is hosted on two servers secured with a corporate Active Directory. The application calls a set of Java-based microservices on an application server and stores data in a clustered MySQL database. The application is heavily used during the day on weekdays. It is lightly used during the evenings and weekends.
Daytime traffic to the application has increased rapidly, and reliability has diminished as a result. The company wants to migrate the application to AWS with a solution that eliminates the need for server maintenance, with an API to securely connect to the microservices.
Which combination of actions will meet these requirements? (Select THREE.)

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

Answer: ACE

NEW QUESTION 117
- (Exam Topic 1)
A financial company is building a system to generate monthly, immutable bank account statements for its users. Statements are stored in Amazon S3. Users should have immediate access to their monthly statements for up to 2 years. Some users access their statements frequently, whereas others rarely access their statements. The company's security and compliance policy requires that the statements be retained for at least 7 years.
What is the MOST cost-effective solution to meet the company's needs?

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

Answer: C

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

NEW QUESTION 119

- (Exam Topic 1)
A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC. A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.
Which solution meets these requirements?

A. Provision a Direct Connect gatewa
B. Delete the existing private virtual interface from the existing connectio
C. Create the second Direct Connect connectio
D. Create a new private virtual interlace on each connection, and connect both private victual interfaces to the Direct Connect gatewa
E. Connect the Direct Connect gateway to the single VPC.
F. Keep the existing private virtual interfac
G. Create the second Direct Connect connectio
H. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC.
I. Keep the existing private virtual interfac
J. Create the second Direct Connect connectio
K. Create a new public virtual interface on the new connection, and connect the new public virtual interface to the single VPC.
L. Provision a transit gatewa
M. Delete the existing private virtual interface from the existing connection.Create the second Direct Connect connectio
N. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gatewa
O. Associate the transit gateway with the single VPC.

Answer: A

Explanation:
A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions. The following describe scenarios where you can use a Direct Connect gateway.
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html

NEW QUESTION 124
- (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 127
- (Exam Topic 1)
A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:
* 1. The data must be highly durable and available.
* 2. The data must always be encrypted at rest and in transit.
* 3. The encryption key must be managed by the company and rotated periodically.
Which of the following solutions should the solutions architect recommend?

A. Deploy the storage gateway to AWS in file gateway mod
B. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.
C. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.
D. Use Amazon DynamoDB with SSL to connect to DynamoD
E. Use an AWS KMS key to encrypt DynamoDB objects at rest.
F. Deploy instances with Amazon EBS volumes attached to store this dat
G. Use E8S volume encryption using an AWS KMS key to encrypt the data.

Answer: B

Explanation:
Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.

NEW QUESTION 131
- (Exam Topic 1)
A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.
What should a solutions architect do to meet these requirements?

A. Create a new developer accoun
B. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organization
C. Enforce a tagging policy that denotes Region affinity.

D. Create an SCP that denies the launch of all EC2 instances except I3.small EC2 instances in us-east-2.Attach the SCP to the project's account.
E. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.
F. Create an IAM policy than allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

Answer: D

NEW QUESTION 135
- (Exam Topic 1)
A solution architect needs to deploy an application on a fleet of Amazon EC2 instances. The EC2 instances run in private subnets in An Auto Scaling group. The application is expected to generate logs at a rate of 100 MB each second on each of the EC2 instances.
The logs must be stored in an Amazon S3 bucket so that an Amazon EMR cluster can consume them for further processing The logs must be quickly accessible for the first 90 days and should be retrievable within 48 hours thereafter.
What is the MOST cost-effective solution that meets these requirements?

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

Answer: C

NEW QUESTION 137
- (Exam Topic 1)
A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance The DB instance is expected to receive many more reads than writes. The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available.
Which steps should the solutions architect take to meet these requirements? (Select THREE)

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

Answer: BCF

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

NEW QUESTION 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 141
- (Exam Topic 1)
A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an
on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.
What should the solutions architect do to meet these requirements?

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

Answer: B

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

NEW QUESTION 143
- (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 146
- (Exam Topic 1)
A company runs a popular public-facing ecommerce website. Its user base is growing quickly from a local
market to a national market. The website is hosted in an on-premises data center with web servers and a MySQL database. The company wants to migrate its workload (o AWS. A solutions architect needs to create a solution to:
• Improve security
• Improve reliability Improve availability
• Reduce latency
• Reduce maintenance
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Use Amazon EC2 instances in two Availability Zones for the web servers in an Auto Scaling group behind an Application Load Balancer.
B. Migrate the database to a Multi-AZ Amazon Aurora MySQL DB cluster.
C. Use Amazon EC2 instances in two Availability Zones to host a highly available MySQL database cluster.
D. Host static website content in Amazon S3. Use S3 Transfer Acceleration to reduce latency while serving webpage
E. Use AWS WAF to improve website security.
F. Host static website content in Amazon S3. Use Amazon CloudFronl to reduce latency while serving webpage
G. Use AWS WAF to improve website security
H. Migrate the database to a single-AZ Amazon RDS for MySQL DB instance.

Answer: ABE

NEW QUESTION 147
- (Exam Topic 1)
A company wants to control its cost of Amazon Athena usage The company has allocated a specific monthly budget for Athena usage A solutions architect must design a solution that will prevent the company from exceeding the budgeted amount
Which solution will moot these requirements?

A. Use AWS Budget
B. Create an alarm (or when the cost of Athena usage reaches the budgeted amount for the mont
C. Configure AWS Budgets actions to deactivate Athena until the end of the month.

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

Answer: D

NEW QUESTION 150
- (Exam Topic 1)
A multimedia company needs to deliver its video-on-demand (VOD) content to its subscribers in a
cost-effective way. The video files range in size from 1-15 GB and are typically viewed frequently for the first 6 months alter creation, and then access decreases considerably. The company requires all video files to remain immediately available for subscribers. There are now roughly 30.000 files, and the company anticipates doubling that number over time.
What is the MOST cost-effective solution for delivering the company's VOD content?

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

Answer: A

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

NEW QUESTION 154
- (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 AmazonS3 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 156
- (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 161
- (Exam Topic 1)
A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a

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

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

Answer: BE

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

NEW QUESTION 163
- (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 167
- (Exam Topic 1)
A company is migrating its three-tier web application from on-premises to the AWS Cloud. The company has the following requirements for the migration process:
• Ingest machine images from the on-premises environment.
• Synchronize changes from the on-premises environment to the AWS environment until the production cutover.
• Minimize downtime when executing the production cutover.
• Migrate the virtual machines' root volumes and data volumes.
Which solution will satisfy these requirements with minimal operational overhead?

A. Use AWS Server Migration Service (SMS) to create and launch a replication job for each tier of the applicatio
B. Launch instances from the AMIs created by AWS SM
C. After initial testing, perform a final replication and create new instances from the updated AMIs.
D. Create an AWS CLIVM Import/Export script to migrate each virtual machin
E. Schedule the script to runincrementally to maintain changes in the applicatio
F. Launch instances from the AMIs created by VM Import/Expor
G. Once testing is done, rerun the script to do a final import and launch the instances from the AMIs.
H. Use AWS Server Migration Service (SMS) to upload the operating system volume
I. Use the AWS CLI import-snaps hot command 'or the data volume
J. Launch instances from the AMIs created by AWS SMS and attach the data volumes to the instance
K. After initial testing, perform a final replication, launch new instances from the replicated AMI
L. and attach the data volumes to the instances.
M. Use AWS Application Discovery Service and AWS Migration Hub to group the virtual machines as an applicatio
N. Use the AWS CLI VM Import/Export script to import the virtual machines as AMI
O. Schedule the script to run incrementally to maintain changes in the applicatio
P. Launch instances from the AMI
Q. After initial testing, perform a final virtual machine import and launch new instances from the AMIs.

Answer: A

Explanation:
SMS can handle migrating the data volumes:
https://aws.amazon.com/about-aws/whats-new/2018/09/aws-server-migration-service-adds-support-for-migratin

NEW QUESTION 172
- (Exam Topic 1)
A company is running an application on Amazon EC2 instances in three environments; development, testing, and production. The company uses AMIs to deploy the EC2 instances. The company builds the AMIs by using custom deployment scripts and infrastructure orchestration tools for each release in each environment. The company is receiving errors in its deployment process. Errors appear during operating system package downloads and during application code installation from a third-party Git hosting service. The company needs deployments to become more reliable across all environments.
Which combination of steps will meet these requirements? (Select THREE).

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

Answer: ACE

NEW QUESTION 173
- (Exam Topic 1)
A start up company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.
The company's existing architecture includes the following:
• A VPC with private and public subnets, and a NAT gateway
• Site-to-Site VPN for connectivity with the on-premises environment
• EC2 security groups with direct SSH access from the on-premises environment
The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers.
Which strategy should a solutions architect use?

A. Install and configure EC2 instance Connect on the fleet of EC2 instance
B. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI.
C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
D. Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.
E. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
F. Enable AWS Config for EC2 security group resource change
G. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.
H. Create an IAM role with the Ama2onSSMManagedInstanceCore managed policy attache
I. Attach the IAM role to all the EC2 instance
J. Remove all security group rules attached to the EC2
K. instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Answer: B

NEW QUESTION 178
- (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 180
- (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

NEW QUESTION 183
- (Exam Topic 1)
A solutions architect at a largo company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization m AWS Organizations The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the interne) The company deploys resources only Into a single AWS Region The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone
Which solution meets these requirements?

A. Creates a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Create an Auto Scaling group of Amazon EC2 Instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Region Modify all default routes to point to the proxy's Auto Scaling group
B. Create a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Use an AWS Network Firewall firewall for rule-based filtering Create Network Firewall endpoints In each Availability Zone Modify all default routes to point to the Network Firewall endpoints
C. Create an AWS Network Firewall firewal for rule-based filtering in each AWS account Modify all default routes to point to the Network Firewall firewalls in each account.
D. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering Modify all default routes to point to the proxy's Auto Scaling group.

Answer: B

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/deployment-models-for-aws-network-firewall/
https://aws.amazon.com/blogs/networking-and-content-delivery/deploy-centralized-traffic-filtering-using-aws-n

NEW QUESTION 186
- (Exam Topic 2)
A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days
The company has a high-speed AWS Direct Connect connection Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day
Which solution meets these requirements?

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

Answer: C

NEW QUESTION 187
.....

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

SAP-C02 Practice Exam Features:

* SAP-C02 Questions and Answers Updated Frequently

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

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

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