Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

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

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

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

Answer: AB

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

NEW QUESTION 2
- (Exam Topic 1)
A company Is serving files to its customers through an SFTP server that Is accessible over the internet The SFTP server Is running on a single Amazon EC2 instance with an Elastic IP address attached Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication The EC2 instance also has an attached security group that allows access from all customer IP addresses.
A solutions architect must implement a solution to improve availability minimize the complexity ot infrastructure management and minimize the disruption to customers who access files. The solution must not change the way customers connect.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 3
- (Exam Topic 1)
A company hosts a large on-premises MySQL database at its main office that supports an issue tracking system used by employees around the world. The company already uses AWS for some workloads and has created an Amazon Route 53 entry for the database endpoint that points to the on-premises database. Management is concerned about the database being a single point of failure and wants a solutions architect to migrate the database to AWS without any data loss or downtime.
Which set of actions should the solutions architect implement?

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

L. This Aurora cluster will be in a master-master replication configuration with the on-premises databas
M. Update the Route 53 entry for the database to point to the Aurora cluster endpoin
N. and shut down the on-premises database.

Answer: C

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

NEW QUESTION 4
- (Exam Topic 1)
A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer The security team requires that all application access attempts be made available for analysis Information about the client IP address, connection type, and user agent must be included.
Which solution will meet these requirements?

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

Answer: C

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

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

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

Answer: C

NEW QUESTION 8
- (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 9
- (Exam Topic 1)
A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.
The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.
Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?

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

Answer: A

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

NEW QUESTION 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 11
- (Exam Topic 1)
A company has developed an application that is running Windows Server on VMware vSphere VMs that the company hosts or premises. The application data is stored in a proprietary format that must be read through the application. The company manually provisioned the servers and the application.
As pan of us disaster recovery plan, the company warns the ability to host its application on AWS temporarily me company's on-premises environment becomes unavailable The company wants the application to return to on-premises hosting after a disaster recovery event is complete The RPO 15 5 minutes.
Which solution meets these requirements with the LEAST amount of operational overhead?

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

Answer: D

NEW QUESTION 15
- (Exam Topic 1)
A company that is developing a mobile game is making game assets available in two AWS Regions. Game assets ate served from a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in each Region. The company requires game assets to be (etched from the closest Region. If game assets become unavailable in the closest Region, they should be fetched from the other Region.
What should a solutions architect do to meet these requirements?

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

Answer: D

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

NEW QUESTION 18
- (Exam Topic 1)
A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.
Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts
Which solution meets these requirements with the LEAST amount of operational overhead?

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

Answer: B

NEW QUESTION 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 22
- (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 26
- (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 30
- (Exam Topic 1)
A solution architect is designing an AWS account structure for a company that consists of multiple terms. All the team will work in the same AWS Region. The company needs a VPC that is connected to the on-premises network. The company expects less than 50 Mbps of total to and from the on-premises network.
Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO)

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

Answer: BD

NEW QUESTION 35
- (Exam Topic 1)
A solutions architect must analyze a company's Amazon EC2 Instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently The company is
running several large, high-memory EC2 instances lo host database dusters that are deployed in active/passive configurations The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern
The solutions architect must analyze the environment and take action based on the findings. Which solution meets these requirements MOST cost-effectively?

A. Create a dashboard by using AWS Systems Manager OpsConter Configure visualizations tor Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes Review the dashboard periodically and identify usage patterns Rightsize the EC2 instances based on the peaks in the metrics
B. Turn on Amazon CloudWatch detailed monitoring for the EC2 instances and their EBS volumes Create and review a dashboard that is based on the metrics Identify usage patterns Rightsize the FC? instances based on the peaks In the metrics

C. Install the Amazon CloudWatch agent on each of the EC2 Instances Turn on AWS Compute Optimizer, and let it run for at least 12 hours Review the recommendations from Compute Optimizer, and rightsize the EC2 instances as directed
D. Sign up for the AWS Enterprise Support plan Turn on AWS Trusted Advisor Wait 12 hours Review the recommendations from Trusted Advisor, and rightsize the EC2 instances as directed

Answer: C

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

NEW QUESTION 38
- (Exam Topic 1)
A company has multiple AWS accounts as part of an organization created with AWS Organizations. Each account has a VPC in the us-east-2 Region and is used for either production or development workloads. Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other.
To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another.
Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

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

Answer: C

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

NEW QUESTION 40
- (Exam Topic 1)
A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.
All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX), and must accommodate high levels of throughput.
Which storage solution will meet these requirements?

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

Answer: D

NEW QUESTION 41
- (Exam Topic 1)
A company is running a web application on Amazon EC2 instances in a production AWS account. The company requires all logs generated from the web application to be copied to a central AWS account (or analysis and archiving. The company's AWS accounts are currently managed independently. Logging agents are configured on the EC2 instances to upload the tog files to an Amazon S3 bucket in the central AWS account.
A solutions architect needs to provide access for a solution that will allow the production account to store log files in the central account. The central account also needs to have read access to the tog files.
What should the solutions architect do to meet these requirements?

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

Answer: B

NEW QUESTION 43
- (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 48
- (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 50
- (Exam Topic 1)
A company that tracks medical devices in hospitals wants to migrate its existing storage solution to the AWS Cloud. The company equips all of its devices with sensors that collect location and usage information. This sensor data is sent in unpredictable patterns with large spikes. The data is stored in a MySQL database running on premises at each hospital. The company wants the cloud storage solution to scale with usage.
The company's analytics team uses the sensor data to calculate usage by device type and hospital. The team needs to keep analysis tools running locally while fetching data from the cloud. The team also needs to use existing Java application and SQL queries with as few changes as possible.
How should a solutions architect meet these requirements while ensuring the sensor data is secure?

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

Answer: C

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

NEW QUESTION 52
- (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 54
- (Exam Topic 1)
A large payroll company recently merged with a small staffing company. The unified company now has multiple business units, each with its own existing AWS account.
A solutions architect must ensure that the company can centrally manage the billing and access policies for all the AWS accounts. The solutions architect configures AWS Organizations by sending an invitation to all member accounts of the company from a centralized management account.
What should the solutions architect do next to meet these requirements?

A. Create the OrganizationAccountAccess IAM group in each member accoun
B. Include the necessary IAM roles for each administrator.
C. Create the OrganizationAccountAccessPolicy IAM policy in each member accoun
D. Connect the member accounts to the management account by using cross-account access.
E. Create the OrganizationAccountAccessRole IAM role in each member accoun
F. Grant permission to the management account to assume the IAM role.
G. Create the OrganizationAccountAccessRole IAM role in the management account Attach the Administrator Access AWS managed policy to the IAM rol
H. Assign the IAM role to the administrators in each member account.

Answer: C

NEW QUESTION 58
- (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 62
- (Exam Topic 1)
A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with the number of files rapidly declining after business hours.
What is the MOST cost-effective migration recommendation?

A. Create a queue using Amazon SQ
B. Configure the existing web server to publish to the new queue.When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
C. Store the processed files in an Amazon S3 bucket.
D. Create a queue using Amazon M
E. Configure the existing web server to publish to the new queue.When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the file
F. Store the processed files in Amazon EF
G. Shut down the EC2 instance after the task is complete.
H. Create a queue using Amazon M
I. Configure the existing web server to publish to the new queue.When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
J. Store the processed files in Amazon EFS.
K. Create a queue using Amazon SO

L. Configure the existing web server to publish to the new queu
M. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the file
N. Scale the EC2 instances based on the SOS queue lengt
O. Store the processed files in an Amazon S3 bucket.

Answer: D

Explanation:
https://aws.amazon.com/blogs/compute/operating-lambda-performance-optimization-part-1/

NEW QUESTION 66
- (Exam Topic 1)
A company needs to store and process image data that will be uploaded from mobile devices using a custom mobile app. Usage peaks between 8 AM and 5 PM on weekdays, with thousands of uploads per minute. The app is rarely used at any other time A user is notified when image processing is complete.
Which combination of actions should a solutions architect take to ensure image processing can scale to handle the load1? (Select THREE.)

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

Answer: BCE

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

NEW QUESTION 71
- (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 75
- (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 79
- (Exam Topic 1)
A company has a photo sharing social networking application. To provide a consistent experience for users, the company performs some image processing on the photos uploaded by users before publishing on the application. The image processing is implemented using a set of Python libraries.
The current architecture is as follows:
• The image processing Python code runs in a single Amazon EC2 instance and stores the processed images in an Amazon S3 bucket named ImageBucket.
• The front-end application, hosted in another bucket, loads the images from ImageBucket to display to users. With plans for global expansion, the company wants to implement changes in its existing architecture to be able to scale for increased demand on the application and reduce management complexity as the application scales.

Which combination of changes should a solutions architect make? (Select TWO.)

A. Place the image processing EC2 instance into an Auto Scaling group.
B. Use AWS Lambda to run the image processing tasks.
C. Use Amazon Rekognition for image processing.
D. Use Amazon CloudFront in front of ImageBucket.
E. Deploy the applications in an Amazon ECS cluster and apply Service Auto Scaling.

Answer: BD

Explanation:
https://prismatic.io/blog/why-we-moved-from-lambda-to-ecs/

NEW QUESTION 84
- (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 85
- (Exam Topic 1)
A large company in Europe plans to migrate its applications to the AWS Cloud. The company uses multiple AWS accounts for various business groups. A data privacy law requires the company to restrict developers'
access to AWS European Regions only.
What should the solutions architect do to meet this requirement with the LEAST amount of management overhead^

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

Answer: B

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

NEW QUESTION 88
- (Exam Topic 1)
A company is building a hybrid solution between its existing on-premises systems and a new backend in AWS. The company has a management application to monitor the state of its current IT infrastructure and automate responses to issues. The company wants to incorporate the status of its consumed AWS services into the application. The application uses an HTTPS endpoint to receive updates.
Which approach meets these requirements with the LEAST amount of operational overhead?

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

Answer: A

Explanation:

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

NEW QUESTION 93
- (Exam Topic 1)
A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes lo the items stored in the DynamoDB tables must be logged within 30 minutes.
Which solution meets the requirements?

A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them (or anomalous behavior
B. Send Amazon SNS notifications when anomalous behaviors are detected.
C. Use AWS CloudTrail to capture all the APIs that change the DynamoDB table
D. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
E. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambd
F. Create a Lambda function to output records lo Amazon Kinesis Data Stream
G. Analyze any anomalies with Amazon Kinesis Data Analytic
H. Send SNS notifications when anomalous behaviors are detected.
I. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda (unction as a target to analyze behavio
J. Send SNS notifications when anomalous behaviors are detected.

Answer: C

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

NEW QUESTION 98
- (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 101
- (Exam Topic 1)
A company plans to migrate to AWS. A solutions architect uses AWS Application Discovery Service over the fleet and discovers that there is an Oracle data warehouse and several PostgreSQL databases. Which combination of migration patterns will reduce licensing costs and operational overhead? (Select TWO.)

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

Answer: BD

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

NEW QUESTION 105
- (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 108
- (Exam Topic 1)
A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.
After initial deployment, the company observes 1.000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.
Which approach should the company take to secure its API?

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

Answer: D

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

NEW QUESTION 109
- (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 111
- (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 114
- (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 117
- (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 118
- (Exam Topic 1)
A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services. Database credentials are hard-coded on each instance. SSH keys for command-tine remote access are stored in a secured Amazon S3 bucket. The company has asked its solutions

architect to improve the security posture of the architecture without adding operational complexity.
Which combination of steps should the solutions architect take to accomplish this? (Select THREE.)

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

Answer: ACF

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

NEW QUESTION 123
- (Exam Topic 1)
A company is using AWS CodePipeline for the CI/CO of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts. As the application has become more complex, recent resource changes in the Cloud Formation templates have caused unplanned downtime.
How should a solutions architect improve the CI'CD pipeline to reduce the likelihood that changes in the templates will cause downtime?

A. Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployment
B. Write test plans for a testing team to execute in a non-production environment before approving the change for production.
C. Implement automated testing using AWS CodeBuild in a test environmen
D. Use CloudFormation changesets to evaluate changes before deploymen
E. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.
F. Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correc
G. Adapt the deployment code to check for error conditions and generate notifications on error
H. Deploy to a test environment and execute a manual test plan before approving the change for production.
I. Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment script
J. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.

Answer: B

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

NEW QUESTION 125
- (Exam Topic 1)
A company has implemented an ordering system using an event-dnven architecture. Dunng initial testing, the system stopped processing orders Further tog analysis revealed that one order message in an Amazon Simple Queue Service (Amazon SOS) standard queue was causing an error on the backend and blocking all subsequent order messages The visibility timeout of the queue is set to 30 seconds, and the backend processing timeout is set to 10 seconds. A solutions architect needs to analyze faulty order messages and ensure that the system continues to process subsequent messages
Which step should the solutions architect take to meet these requirements?

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

Answer: D

NEW QUESTION 126
- (Exam Topic 1)
A company runs a popular web application in an on-premises data center. The application receives four million views weekly. The company expects traffic to increase by 200% because of an advertisement that will be published soon.
The company needs to decrease the load on the origin before the increase of traffic occurs. The company does not have enough time to move the entire application to the AWS Cloud.
Which solution will meet these requirements?

A. Create an Amazon CloudFront content delivery network (CDN). Enable query forwarding to the origin.Create a managed cache policy that includes query string
B. Use an on-premises load balancer as the origi
C. Offload the DNS querying to AWS to handle CloudFront CDN traffic.
D. Create an Amazon CloudFront content delivery network (CDN) that uses a Real Time Messaging Protocol (RTMP) distributio
E. Enable query forwarding to the origi
F. Use an on-premises load balancer as the origi
G. Offload the DNS querying to AWS to handle CloudFront CDN traffic.
H. Create an accelerator in AWS Global Accelerato
I. Add listeners for HTTP and HTTPS TCP ports.Create an endpoint grou
J. Create a Network Load Balancer (NLB), and attach it to the endpoint grou
K. Point the NLB to the on-premises server
L. Offload the DNS querying to AWS to handle AWS Global Accelerator traffic.
M. Create an accelerator in AWS Global Accelerato
N. Add listeners for HTTP and HTTPS TCP ports.Create an endpoint grou
O. Create an Application Load Balancer (ALB), and attach it to the endpoint grou
P. Point the ALB to the on-premises server
Q. Offload the DNS querying to AWS to handle AWS Global Accelerator traffic.

Answer: D

NEW QUESTION 128
- (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 129
- (Exam Topic 1)
A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.
The company wants to stop using email in the application and move to a serverless application model. Which architecture solution meets these requirements?

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

Answer: C

NEW QUESTION 133
- (Exam Topic 1)
An ecommerce website running on AWS uses an Amazon RDS for MySQL DB instance with General Purpose SSD storage. The developers chose an appropriate instance type based on demand, and configured 100 GB of storage with a sufficient amount of free space.
The website was running smoothly for a few weeks until a marketing campaign launched. On the second day of the campaign, users reported long wait times and time outs. Amazon CloudWatch metrics indicated that both reads and writes to the DB instance were experiencing long response times. The CloudWatch metrics show 40% to 50% CPU and memory utilization, and sufficient free storage space is still available. The application server logs show no evidence of database connectivity issues.
What could be the root cause of the issue with the marketing campaign?

A. It exhausted the I/O credit balance due to provisioning low disk storage during the setup phase.
B. It caused the data in the tables to change frequently, requiring indexes to be rebuilt to optimize queries.
C. It exhausted the maximum number of allowed connections to the database instance.
D. It exhausted the network bandwidth available to the RDS for MySQL DB instance.

Answer: A

Explanation:
"When using General Purpose SSD storage, your DB instance receives an initial I/O credit balance of 5.4 million I/O credits. This initial credit balance is enough to sustain a burst performance of 3,000 IOPS for 30 minutes."
https://aws.amazon.com/blogs/database/how-to-use-cloudwatch-metrics-to-decide-between-general-purpose-or

NEW QUESTION 134
- (Exam Topic 2)
A finance company is storing financial records in an Amazon S3 bucket. The company persists a record for every financial transaction. According to regulatory requirements, the records cannot be modified for at least 1 year after they are written. The records are read on a regular basis and must be immediately accessible.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 135
- (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 137
- (Exam Topic 2)
A company is migrating its marketing website and content management system from an on-premises data center to AWS. The company wants the AWS application to be deployed in a VPC with Amazon EC2 instances used for the web servers and an Amazon RDS instance for the database.
The company has a runbook document that describes the installation process of the on-premises system. The company would like to base the AWS system on the processes referenced in the runbook document. The runbook document describes the installation and configuration of the operating systems, network settings, the website, and content management system software on the servers After the migration is complete, the company wants to be able to make changes quickly to take advantage of other AWS features.
How can the application and environment be deployed and automated m AWS. while allowing for future changes?

A. Update the runbook to describe how to create the VP
B. the EC2 instances and the RDS instance for the application by using the AWS Console Make sure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration
C. Write a Python script that uses the AWS API to create the VP
D. the EC2 instances and the RDS instance for the application Write shell scripts that implement the rest of the steps in the runbook Have the Python script copy and run the shell scripts on the newly created instances to complete the installation
E. Write an AWS Cloud Formation template that creates the VPC, the EC2 instances, and the RDS instance for the application Ensure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration
F. Write an AWS CloudFormation template that creates the VPC the EC2 instances, and the RDS instance for the application Include EC2 user data in the AWS Cloud Formation template to install and configure the software.

Answer: D

NEW QUESTION 141
- (Exam Topic 2)
A company wants to allow its marketing team to perform SQL queries on customer records to identify market segments. The data is spread across hundreds of files. The records must be encrypted in transit and at rest. The team manager must have the ability to manage users and groups but no team members should have access to services or resources not required for the SQL queries Additionally, administrators need to audit the queries made and receive notifications when a query violates rules defined by the security team.
AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the team manager. Which design meets these requirements'?

A. Apply a service control policy (SCP) that allows access to IAM Amazon RD
B. and AWS CloudTrail Load customer records in Amazon RDS MySQL and train users to run queries using the AWS CL
C. Stream the query logs to Amazon CloudWatch Logs from the RDS database instance Use a subscription filter with AWS Lambda functions to audit and alarm on queries against personal data
D. Apply a service control policy (SCP) that denies access to all services except IAM Amazon Athena Amazon S3 and AWS CloudTrail Store customer record files in Amazon S3 and tram users to run queries using the CLI via Athena Analyze CloudTrail events to audit and alarm on queries against personal data
E. Apply a service control policy (SCP) that denies access to all services except IAM Amazon DynamoD
F. and AWS CloudTrail Store customer records in DynamoDB and train users to run queries using the AWS CLI Enable DynamoDB streams to track the queries that are issued and use an AWS Lambda function for real-time monitoring and alerting
G. Apply a service control policy (SCP) that allows access to IAM Amazon Athena; Amazon S3, and AWS CloudTrail Store customer records as files in Amazon S3 and train users to leverage the Amazon S3 Select feature and run queries using the AWS CLI Enable S3 object-level logging and analyze CloudTrail events to audit and alarm on queries against personal data

Answer: B

NEW QUESTION 143
- (Exam Topic 2)
A company runs a proprietary stateless ETL application on an Amazon EC2 Linux instance. The application is a Linux binary, and the source code cannot be modified. The application is single-threaded, uses 2 GB of RAM. and is highly CPU intensive The application is scheduled to run every 4 hours and runs for up to 20 minutes A solutions architect wants to revise the architecture for the solution.
Which strategy should the solutions architect use?

A. Use AWS Lambda to run the applicatio

B. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours
C. Use AWS Batch to run the application Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours
D. Use AWS Fargate to run the application Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours
E. Use Amazon 6C2 Spot Instances to run the application Use AWS CodeDeptoy to deploy and run the application every 4 hours.

Answer: C

NEW QUESTION 144
- (Exam Topic 2)
A company wants to migrate its website from an on-premises data center onto AWS At the same time it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege
A solutions architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster What steps are required after the deployment to meet the requirements'? (Select TWO.)

A. Create tasks using the bridge network mode
B. Create tasks using the awsvpc network mode
C. Apply security groups to Amazon EC2 instances and use IAM roles for EC2 instances to access other resources
D. Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources
E. Apply security groups to the tasks; and use IAM roles for tasks to access other resources

Answer: BE

NEW QUESTION 146
- (Exam Topic 2)
A company is running a three-tier web application in an on-premises data center. The frontend is served by an Apache web server, the middle tier is a monolithic Java application, and the storage tier is a PostgreSOL database.
During a recent marketing promotion, customers could not place orders through the application because the application crashed An analysis showed that all three tiers were overloaded. The application became unresponsive, and the database reached its capacity limit because of read operations. The company already has several similar promotions scheduled in the near future.
A solutions architect must develop a plan for migration to AWS to resolve these issues. The solution must maximize scalability and must minimize operational effort.
Which combination of steps will meet these requirements? (Select THREE.)

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

Answer: BCF

NEW QUESTION 148
- (Exam Topic 2)
A fleet of Amazon ECS instances is used to poll an Amazon SQS queue and update items in an Amazon DynamoDB database Items in the table are not being updated, and the SQS queue Is filling up Amazon CloudWatch Logs are showing consistent 400 errors when attempting to update the table The provisioned write capacity units are appropriately configured, and no throttling is occurring
What is the LIKELY cause of the failure*?

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

Answer: D

NEW QUESTION 153
- (Exam Topic 2)
A company has a new security policy. The policy requires the company to log any event that retrieves data from Amazon S3 buckets. The company must save these audit logs in a dedicated S3 bucket. The company created the audit logs S3 bucket in an AWS account that is designated for centralized logging. The S3 bucket has a bucket policy that allows write-only cross-account access A solutions architect must ensure that all S3 object-level access is being logged for current S3 buckets and future S3 buckets. Which solution will meet these requirements?

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

Answer: D

NEW QUESTION 158
- (Exam Topic 2)
A company uses multiple AWS accounts in a single AWS Region A solutions architect is designing a solution to consolidate logs generated by Elastic Load Balancers (ELBs) in the AppDev, AppTest and AppProd accounts. The logs should be stored in an existing Amazon S3 bucket named s3-eib-logs in the central AWS account. The central account is used for log consolidation only and does not have ELBs deployed ELB logs must be encrypted at rest
Which combination of steps should the solutions architect take to build the solution" (Select TWO)

A. Update the S3 bucket policy for the s3-elb-logs bucket to allow the s3 PutBucketLogging action for the central AWS account ID
B. Update the S3 bucket policy for the s3-eib-logs bucket to allow the s3 PutObject and s3 DeleteObject actions for the AppDev AppTest and AppProd account IDs
C. Update the S3 bucket policy for the s3-elb-logs bucket to allow the s3 PutObject action for the AppDev AppTest and AppProd account IDs
D. Enable access logging for the ELB
E. Set the S3 location to the s3-elb-logs bucket
F. Enable Amazon S3 default encryption using server-side encryption with S3 managed encryption keys (SSE-S3) for the s3-elb-logs S3 bucket

Answer: AE

NEW QUESTION 163
- (Exam Topic 2)
A company is planning to migrate an application from on premises to the AWS Cloud. The company will begin the migration by moving the application's underlying data storage to AWS The application data is stored on a shared tie system on premises, and the application servers connect to the shared We system through SMB.
A solutions architect must implement a solution that uses an Amazon S3 bucket tor shared storage Until the application Is fully migrated and code is rewritten to use native Amazon S3 APIs, the application must continue to have access to the data through SMB The solutions architect must migrate the application data to AWS to its new location while still allowing the on-premises application to access the data.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 166
- (Exam Topic 2)
A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts AWS Site-to-Site VPN connections are configured between ail of the company's global offices and the transit account The company has AWS Config enabled on all of its accounts.
The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices Developers Will reference this list to gain access to applications securely.
Which solution meets these requirements with the LEAST amount of operational overhead?

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

Answer: C

NEW QUESTION 169
- (Exam Topic 2)
A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.
The website contains stat c content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2 instances running in an Auto Scaling group to process an Amazon SQS queue The company wants to re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.
Which solution meets these requirements?

A. Use Amazon ECS containers for the web application and Spot Instances for the Auto Scaling group that processes the SQS queu
B. Replace the custom software with Amazon Recognition to categorize the videos.
C. Store the uploaded videos n Amazon EFS and mount the file system to the EC2 instances for Te web applicatio
D. Process the SOS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
E. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notifications to publish events to the SQS queue Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
F. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the web application and launch a worker environment to process the SQS

queue Replace the custom software with Amazon Rekognition to categorize the videos.

Answer: D

NEW QUESTION 174
- (Exam Topic 2)
A company is using a lift-and-shift strategy to migrate applications from several on-premises Windows servers to AWS. The Windows servers will be hosted on Amazon EC2 instances in the us-east-1 Region.
The company's security policy allows the installation of migration tools on servers. The migration data must be encrypted in transit and encrypted at rest. The applications are business critical. The company needs to minimize the cutover window and minimize the downtime that results from the migration. The company wants to use Amazon CloudWatch and AWS CloudTrail for monitoring.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 176
- (Exam Topic 2)
A company is migrating its data centre from on premises to the AWS Cloud. The migration will take several months to complete. The company will use Amazon Route 53 for private DNS zones.
During the migration, the company must Keep its AWS services pointed at the VPC's Route 53 Resolver for DNS. The company also must maintain the ability to resolve addresses from its on-premises DNS server A solutions architect must set up DNS so that Amazon EC2 instances can use native Route 53 endpoints to resolve on-premises DNS queries
Which configuration writ meet these requirements?

A. Configure Vie VPC DHCP options set to point to on-premises DNS server IP addresse
B. Ensure that security groups for EC2 instances allow outbound access to port 53 on those DNS server IP addresses.
C. Launch an EC2 instance that has DNS BIND installed and configure
D. Ensure that the security groups that are attached to the EC2 instance can access the on-premises DNS server IP address on port 53. Configure BIND to forward DNS queries to on-premises DNS server IP addresses Configure each migrated EC2 instances DNS settings to point to the BIND server IP address.
E. Create a new outbound endpoint in Route 53. and attach me endpoint to the VP
F. Ensure that the security groups that are attached to the endpoint can access the on-premises DNS server IP address on port 53 Create a new Route 53 Resolver rule that routes on-premises designated traffic to theon-premises DNS server.
G. Create a new private DNS zone in Route 53 with the same domain name as the on-premises domain.Create a single wildcard record with the on-premises DNS server IP address as the record's address.

Answer: A

NEW QUESTION 181
- (Exam Topic 2)
A solutions architect has been assigned to migrate a 50 TB Oracle data warehouse that contains sales data from on-premises to Amazon Redshift Major updates to the sales data occur on the final calendar day of the month For the remainder of the month, the data warehouse only receives minor daily updates and is primarily used for reading and reporting Because of this the migration process must start on the first day of the month and must be complete before the next set of updates occur. This provides approximately 30 days to complete the migration and ensure that the minor daily changes have been synchronized with the Amazon Redshift data warehouse Because the migration cannot impact normal business network operations, the bandwidth allocated to the migration for moving data over the internet is 50 Mbps The company wants to keep data migration costs low
Which steps will allow the solutions architect to perform the migration within the specified timeline?

A. Install Oracle database software on an Amazon EC2 instance Configure VPN connectivity between AWS and the company's data center Configure the Oracle database running on Amazon EC2 to join the Oracle Real Application Clusters (RAC) When the Oracle database on Amazon EC2 finishes synchronizing, create an AWS DMS ongoing replication task to migrate the data from the Oracle database on Amazon EC2 to Amazon Redshift Verify the data migration is complete and perform the cut over to Amazon Redshift.
B. Create an AWS Snowball import job Export a backup of the Oracle data warehouse Copy the exported data to the Snowball device Return the Snowball device to AWS Create an Amazon RDS for Oracle database and restore the backup file to that RDS instance Create an AWS DMS task to migrate the data from the RDS for Oracle database to Amazon Redshift Copy daily incremental backups from Oracle in the data center to the RDS for Oracle database over the internet Verify the data migration is complete and perform the cut over to Amazon Redshift.
C. Install Oracle database software on an Amazon EC2 instance To minimize the migration time configure VPN connectivity between AWS and the company's data center by provisioning a 1 Gbps AWS Direct Connect connection Configure the Oracle database running on Amazon EC2 to be a read replica of the data center Oracle database Start the synchronization process between the company's on-premises data center and the Oracle database on Amazon EC2 When the Oracle database on Amazon EC2 is synchronized with the on-premises database create an AWS DMS ongoing replication task from the Oracle database read replica that is running on Amazon EC2 to Amazon Redshift Verify the data migration is complete and perform the cut over to Amazon Redshift.
D. Create an AWS Snowball import jo
E. Configure a server in the company€™s data center with an extraction agen
F. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schem
G. Create a new project in AWS SCT using the registered data extraction agen
H. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing change
I. Copy data to the Snowball device and return the Snowball device to AW

J. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshif
K. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

Answer: D

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

NEW QUESTION 182
- (Exam Topic 2)
A company has an on-premises Microsoft SQL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS.
Which solution meets these requirements Most cost effectively?

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

Answer: A

NEW QUESTION 186
- (Exam Topic 2)
A software development company has multiple engineers who are working remotely. The company is running Active Directory Domain Services (AD DS) on an Amazon EC2 instance. The company's security policy states that all internal, nonpublic services that are deployed in a VPC must be accessible through a VPN Multi-factor authentication (MFA) must be used for access to a VPN.
Whet should a solution architect do to meet these requirements?

A. Create an AWS Site-to-Site VPN connection Configure integration between a VPN and AD D
B. Use an Amazon Workspaces client with MFA support enabled to establish a VPN connection.
C. Create an AWS Client VPN endpoint Create an AD Connector directory for integration with AD DS Enable MFA for AD Connector Use AWS Client VPN to establish a VPN connection.
D. Create multiple AWS Site-to-Site VPN connections by using AWS VPN CloudHub Configure integration between AWS VPN CloudHub and AD DS Use AWS Cop4ot to establish a VPN connection.
E. Create an Amazon WorkLink endpoint Configure integration between Amazon WorkLink and AD D
F. Enable MFA in Amazon WorkLink Use AWS Client VPN to establish a VPN connection.

Answer: B

NEW QUESTION 187
- (Exam Topic 2)
A company is migrating its infrastructure to the AW5 Cloud. The company must comply with a variety of regulatory standards for different projects. The company needs a multi-account environment.
A solutions architect needs to prepare the baseline infrastructure The solution must provide a consistent baseline of management and security but it must allow flexibility for different compliance requirements within various AWS accounts. The solution also needs to integrate with the existing on-premises Active Directory Federation Services (AD FS) server.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Create an organization In AWS Organizations Create a single SCP for least privilege access across all accounts Create a single OU for all accounts Configure an IAM identity provider tor federation with the on-premises AD FS server Configure a central togging account with a defined process for log generating services to send log events to the central accoun
B. Enable AWS Config in the central account with conformance packs for all accounts.
C. Create an organization In AWS Organizations Enable AWS Control Tower on the organizatio
D. Review included guardrails for SCP
E. Check AWS Config for areas that require additions Add OUs as necessary Connect AWS Single Sign-On to the on-premises AD FS server
F. Create an organization in AWS Organizations Create SCPs for least privilege access Create an OU structure, and use it to group AWS accounts Connect AWS Single Sign-On to the on-premises AD FS serve
G. Configure a central logging account with a defined process for tog generating services to send log events to the central account Enable AWS Config in the central account with aggregators and conformance packs.
H. Create an organization in AWS Organizations Enable AWS Control Tower on the organization Review included guardrails for SCP
I. Check AWS Config for areas that require additions Configure an IAM identity provider for federation with the on-premises AD FS server.

Answer: A

NEW QUESTION 189

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

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

Answer: A

NEW QUESTION 190
- (Exam Topic 2)
A large company recently experienced an unexpected increase in Amazon RDS and Amazon DynamoDB costs The company needs to increase visibility into details of AWS Billing and Cost Management There are various accounts associated with AWS Organizations, including many development and production accounts. There is no consistent tagging strategy across the organization, but there are guidelines in place that require all infrastructure to be deployed using AWS Cloud Formation with consistent tagging Management requires cost center numbers and project ID numbers for all existing and future DynamoDB tables and RDS instances
Which strategy should the solutions architect provide to meet these requirements?

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

Answer: B

NEW QUESTION 194
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

SAP-C02 Practice Exam Features:

* SAP-C02 Questions and Answers Updated Frequently

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

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

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