Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/

NEW QUESTION 1
- (Exam Topic 1)
A solutions architect is designing the cloud architecture for a new application being deployed on AWS. The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed. The processor application is stateless. The solutions architect must ensure that the application is loosely coupled and the job items are durably stored.
Which design should the solutions architect use?

A. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch configuration that uses the AMI Create an Auto Scaling group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on CPU usage
B. Create an Amazon SQS queue to hold the jobs that need to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch configuration that uses the AM' Create an Auto Scaling group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on network usage
C. Create an Amazon SQS queue to hold the jobs that needs to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue
D. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of messages published to the SNS topic

Answer: C

Explanation:
"Create an Amazon SQS queue to hold the jobs that needs to be processed. Create an Amazon EC2 Auto Scaling group for the compute application. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue"
In this case we need to find a durable and loosely coupled solution for storing jobs. Amazon SQS is ideal for this use case and can be configured to use dynamic scaling based on the number of jobs waiting in the queue.To configure this scaling you can use the backlog per instance metric with the target value being the acceptable backlog per instance to maintain. You can calculate these numbers as follows: Backlog per instance: To calculate your backlog per instance, start with the ApproximateNumberOfMessages queue attribute to determine the length of the SQS queue

NEW QUESTION 2
- (Exam Topic 1)
A company has a production workload that runs on 1,000 Amazon EC2 Linux instances. The workload is powered by third-party software. The company needs to patch the third-party software on all EC2 instances as quickly as possible to remediate a critical security vulnerability.
What should a solutions architect do to meet these requirements?

A. Create an AWS Lambda function to apply the patch to all EC2 instances.
B. Configure AWS Systems Manager Patch Manager to apply the patch to all EC2 instances.
C. Schedule an AWS Systems Manager maintenance window to apply the patch to all EC2 instances.
D. Use AWS Systems Manager Run Command to run a custom command that applies the patch to all EC2 instances.

Answer: B

Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/about-windows-app-patching.html

NEW QUESTION 3
- (Exam Topic 1)
A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a high-speed internet connection. The company's weather forecasting applications are based in a single Region and analyze the data daily.
What is the FASTEST way to aggregate data from all of these global sites?

A. Enable Amazon S3 Transfer Acceleration on the destination bucke
B. Use multipart uploads to directly upload site data to the destination bucket.
C. Upload site data to an Amazon S3 bucket in the closest AWS Regio
D. Use S3 cross-Region replication to copy objects to the destination bucket.
E. Schedule AWS Snowball jobs daily to transfer data to the closest AWS Regio
F. Use S3 cross-Region replication to copy objects to the destination bucket.
G. Upload the data to an Amazon EC2 instance in the closest Regio
H. Store the data in an Amazon Elastic Block Store (Amazon EBS) volum
I. Once a day take an EBS snapshot and copy it to the centralized Regio
J. Restore the EBS volume in the centralized Region and run an analysis on the data daily.

Answer: A

Explanation:
You might want to use Transfer Acceleration on a bucket for various reasons, including the following: You have customers that upload to a centralized bucket from all over the world.
You transfer gigabytes to terabytes of data on a regular basis across continents.
You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3.
https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20redu "Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50-500% for long-distance transfer of larger objects. Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet"
https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html "Improved throughput - You can upload parts in parallel to improve throughput."

NEW QUESTION 4

- (Exam Topic 1)
A company performs monthly maintenance on its AWS infrastructure. During these maintenance activities, the company needs to rotate the credentials tor its Amazon ROS tor MySQL databases across multiple AWS Regions
Which solution will meet these requirements with the LEAST operational overhead?

A. Store the credentials as secrets in AWS Secrets Manage
B. Use multi-Region secret replication for the required Regions Configure Secrets Manager to rotate the secrets on a schedule
C. Store the credentials as secrets in AWS Systems Manager by creating a secure string parameter Use multi-Region secret replication for the required Regions Configure Systems Manager to rotate the secrets on a schedule
D. Store the credentials in an Amazon S3 bucket that has server-side encryption (SSE) enabled Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function to rotate the credentials
E. Encrypt the credentials as secrets by using AWS Key Management Service (AWS KMS) multi-Region customer managed keys Store the secrets in an Amazon DynamoDB global table Use an AWS Lambda function to retrieve the secrets from DynamoDB Use the RDS API to rotate the secrets.

Answer: A

Explanation:
https://aws.amazon.com/blogs/security/how-to-replicate-secrets-aws-secrets-manager-multiple-regions/

NEW QUESTION 5
- (Exam Topic 1)
A company is building an ecommerce web application on AWS. The application sends information about new orders to an Amazon API Gateway REST API to process. The company wants to ensure that orders are processed in the order that they are received.
Which solution will meet these requirements?

A. Use an API Gateway integration to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when the application receives an orde
B. Subscribe an AWS Lambda function to the topic to perform processing.
C. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) FIFO queue when the application receives an orde
D. Configure the SQS FIFO queue to invoke an AWS Lambda function for processing.
E. Use an API Gateway authorizer to block any requests while the application processes an order.
F. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) standard queue when the application receives an orde
G. Configure the SQS standard queue to invoke an AWS Lambda function for processing.

Answer: B

NEW QUESTION 6
- (Exam Topic 1)
A company is preparing to store confidential data in Amazon S3 For compliance reasons the data must be encrypted at rest Encryption key usage must be logged tor auditing purposes. Keys must be rotated every year.
Which solution meets these requirements and «the MOST operationally efferent?

A. Server-side encryption with customer-provided keys (SSE-C)
B. Server-side encryption with Amazon S3 managed keys (SSE-S3)
C. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with manual rotation
D. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with automate rotation

Answer: D

Explanation:
https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html
When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted.
Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use. AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by default on customer managed CMKs. When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

NEW QUESTION 7
- (Exam Topic 1)
A solutions architect is designing a new hybrid architecture to extend a company s on-premises infrastructure to AWS The company requires a highly available connection with consistent low latency to an AWS Region. The company needs to minimize costs and is willing to accept slower traffic if the primary connection fails.
What should the solutions architect do to meet these requirements?

A. Provision an AWS Direct Connect connection to a Region Provision a VPN connection as a backup if the primary Direct Connect connection fails.
B. Provision a VPN tunnel connection to a Region for private connectivit
C. Provision a second VPN tunnel for private connectivity and as a backup if the primary VPN connection fails.
D. Provision an AWS Direct Connect connection to a Region Provision a second Direct Connect connection to the same Region as a backup if the primary Direct Connect connection fails.
E. Provision an AWS Direct Connect connection to a Region Use the Direct Connect failover attribute from the AWS CLI to automatically create a backup connection if the primary Direct Connect connection fails.

Answer: A

Explanation:
"In some cases, this connection alone is not enough. It is always better to guarantee a fallback connection as the backup of DX. There are several options, but implementing it with an AWS Site-To-Site VPN is a real cost-effective solution that can be exploited to reduce costs or, in the meantime, wait for the setup of a second DX."
https://www.proud2becloud.com/hybrid-cloud-networking-backup-aws-direct-connect-network-connection-with

NEW QUESTION 8
- (Exam Topic 1)
A company runs a shopping application that uses Amazon DynamoDB to store customer information. In case of data corruption, a solutions architect needs to design a solution that meets a recovery point objective (RPO) of 15 minutes and a recovery time objective (RTO) of 1 hour.
What should the solutions architect recommend to meet these requirements?

A. Configure DynamoDB global table
B. For RPO recovery, point the application to a different AWS Region.
C. Configure DynamoDB point-in-time recover
D. For RPO recovery, restore to the desired point in time.
E. Export the DynamoDB data to Amazon S3 Glacier on a daily basi
F. For RPO recovery, import the data from S3 Glacier to DynamoDB.
G. Schedule Amazon Elastic Block Store (Amazon EBS) snapshots for the DynamoDB table every 15 minute
H. For RPO recovery, restore the DynamoDB table by using the EBS snapshot.

Answer: B

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/PointInTimeRecovery.html

NEW QUESTION 9
- (Exam Topic 1)
A company is hosting a web application on AWS using a single Amazon EC2 instance that stores
user-uploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone placing both behind an Application Load Balancer After completing this change, users reported that, each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time.
What should a solutions architect propose to ensure users see all of their documents at once?

A. Copy the data so both EBS volumes contain all the documents.
B. Configure the Application Load Balancer to direct a user to the server with the documents
C. Copy the data from both EBS volumes to Amazon EFS Modify the application to save new documents to Amazon EFS
D. Configure the Application Load Balancer to send the request to both servers Return each document from the correct server.

Answer: A

Explanation:
Amazon EFS provides file storage in the AWS Cloud. With Amazon EFS, you can create a file system, mount the file system on an Amazon EC2 instance, and then read and write data to and from your file system. You can mount an Amazon EFS file system in your VPC, through the Network File System versions 4.0 and a4.1 (NFSv4) protocol. We recommend using a current generation Linux NFSv4.1 client, such as those found in the latest Amazon Linux, Redhat, and Ubuntu AMIs, in conjunction with the Amazon EFS Mount Helper. For instructions, see Using the amazon-efs-utils Tools.
For a list of Amazon EC2 Linux Amazon Machine Images (AMIs) that support this protocol, see NFS Support. For some AMIs, you'll need to install an NFS client to mount your file system on your Amazon EC2 instance. For instructions, see Installing the NFS Client.
You can access your Amazon EFS file system concurrently from multiple NFS clients, so applications that scale beyond a single connection can access a file system. Amazon EC2 instances running in multiple Availability Zones within the same AWS Region can access the file system, so that many users can access and share a common data source.
https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-ec2

NEW QUESTION 10
- (Exam Topic 1)
A company has an automobile sales website that stores its listings in a database on Amazon RDS When an automobile is sold the listing needs to be removed from the website and the data must be sent to multiple target systems.
Which design should a solutions architect recommend?

A. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS> queue for the targets to consume
B. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS) FIFO queue for the targets to consume
C. Subscribe to an RDS event notification and send an Amazon Simple Queue Service (Amazon SQS) queue fanned out to multiple Amazon Simple Notification Service (Amazon SNS) topics Use AWS Lambda functions to update the targets
D. Subscribe to an RDS event notification and send an Amazon Simple Notification Service (Amazon SNS) topic fanned out to multiple Amazon Simple Queue Service (Amazon SQS) queues Use AWS Lambda functions to update the targets

Answer: D

Explanation:
https://docs.aws.amazon.com/lambda/latest/dg/services-rds.html https://docs.aws.amazon.com/lambda/latest/dg/with-sns.html

NEW QUESTION 10
- (Exam Topic 1)
A company hosts its multi-tier applications on AWS. For compliance, governance, auditing, and security, the company must track configuration changes on its AWS resources and record a history of API calls made to these resources.
What should a solutions architect do to meet these requirements?

A. Use AWS CloudTrail to track configuration changes and AWS Config to record API calls
B. Use AWS Config to track configuration changes and AWS CloudTrail to record API calls
C. Use AWS Config to track configuration changes and Amazon CloudWatch to record API calls
D. Use AWS CloudTrail to track configuration changes and Amazon CloudWatch to record API calls

Answer: B

NEW QUESTION 14
- (Exam Topic 1)
A company recently signed a contract with an AWS Managed Service Provider (MSP) Partner for help with an application migration initiative. A solutions architect needs to share an Amazon Machine Image (AMI) from an existing AWS account with the MSP Partner's AWS account. The AMI is backed by Amazon Elastic Block Store (Amazon EBS) and uses a customer managed customer master key (CMK) to encrypt EBS volume snapshots.
What is the MOST secure way for the solutions architect to share the AMI with the MSP Partner's AWS account?

A. Make the encrypted AMI and snapshots publicly availabl
B. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key
C. Modify the launchPermission property of the AM
D. Share the AMI with the MSP Partner's AWS account onl
E. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key.
F. Modify the launchPermission property of the AMI Share the AMI with the MSP Partner's AWS account onl
G. Modify the CMK's key policy to trust a new CMK that is owned by the MSP Partner for encryption.
H. Export the AMI from the source account to an Amazon S3 bucket in the MSP Partner's AWS account.Encrypt the S3 bucket with a CMK that is owned by the MSP Partner Copy and launch the AMI in the MSP Partner's AWS account.

Answer: B

Explanation:
Share the existing KMS key with the MSP external account because it has already been used to encrypt the AMI snapshot.
https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external-accounts.html

NEW QUESTION 18
- (Exam Topic 1)
A company wants to improve its ability to clone large amounts of production data into a test environment in the same AWS Region. The data is stored in Amazon EC2 instances on Amazon Elastic Block Store (Amazon EBS) volumes. Modifications to the cloned data must not affect the production environment. The software that accesses this data requires consistently high I/O performance.
A solutions architect needs to minimize the time that is required to clone the production data into the test environment.
Which solution will meet these requirements?

A. Take EBS snapshots of the production EBS volume
B. Restore the snapshots onto EC2 instance store volumes in the test environment.
C. Configure the production EBS volumes to use the EBS Multi-Attach featur
D. Take EBS snapshots of the production EBS volume
E. Attach the production EBS volumes to the EC2 instances in the test environment.
F. Take EBS snapshots of the production EBS volume
G. Create and initialize new EBS volume
H. Attach the new EBS volumes to EC2 instances in the test environment before restoring the volumes from the production EBS snapshots.
I. Take EBS snapshots of the production EBS volume
J. Turn on the EBS fast snapshot restore feature on the EBS snapshot
K. Restore the snapshots into new EBS volume
L. Attach the new EBS volumes to EC2 instances in the test environment.

Answer: C

NEW QUESTION 21
- (Exam Topic 1)
A company is storing backup files by using Amazon S3 Standard storage. The files are accessed frequently for 1 month. However, the files are not accessed after 1 month. The company must keep the files indefinitely.
Which storage solution will meet these requirements MOST cost-effectively?

A. Configure S3 Intelligent-Tiering to automatically migrate objects.
B. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month.
C. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) after 1 month.
D. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 1 month.

Answer: B

NEW QUESTION 23
- (Exam Topic 1)
A company wants to reduce the cost of its existing three-tier web architecture. The web, application, and database servers are running on Amazon EC2 instances for the development, test, and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours.
The production EC2 instances run 24 hours a day. The development and test EC2 instances run for at least 8 hours each day. The company plans to implement automation to stop the development and test EC2 instances when they are not in use.
Which EC2 instance purchasing solution will meet the company's requirements MOST cost-effectively?

A. Use Spot Instances for the production EC2 instance
B. Use Reserved Instances for the development and test EC2 instances.
C. Use Reserved Instances for the production EC2 instance
D. Use On-Demand Instances for the development and test EC2 instances.
E. Use Spot blocks for the production EC2 instance
F. Use Reserved Instances for the development and test EC2 instances.
G. Use On-Demand Instances for the production EC2 instance
H. Use Spot blocks for the development and test EC2 instances.

Answer: B

NEW QUESTION 24
- (Exam Topic 1)
A company's HTTP application is behind a Network Load Balancer (NLB). The NLB's target group is configured to use an Amazon EC2 Auto Scaling group with multiple EC2 instances that run the web service.
The company notices that the NLB is not detecting HTTP errors for the application. These errors require a manual restart of the EC2 instances that run the web service. The company needs to improve the application's availability without writing custom scripts or code.
What should a solutions architect do to meet these requirements?

A. Enable HTTP health checks on the NL
B. supplying the URL of the company's application.
C. Add a cron job to the EC2 instances to check the local application's logs once each minut
D. If HTTP errors are detected, the application will restart.
E. Replace the NLB with an Application Load Balance
F. Enable HTTP health checks by supplying the URL of the company's applicatio
G. Configure an Auto Scaling action to replace unhealthy instances.
H. Create an Amazon Cloud Watch alarm that monitors the UnhealthyHostCount metric for the NL
I. Configure an Auto Scaling action to replace unhealthy instances when the alarm is in the ALARM state.

Answer: C

NEW QUESTION 26
- (Exam Topic 1)
A development team needs to host a website that will be accessed by other teams. The website contents consist of HTML, CSS, client-side JavaScript, and images Which method is the MOST cost-effective for hosting the website?

A. Containerize the website and host it in AWS Fargate.
B. Create an Amazon S3 bucket and host the website there
C. Deploy a web server on an Amazon EC2 instance to host the website.
D. Configure an Application Loa d Balancer with an AWS Lambda target that uses the Express js framework.

Answer: B

Explanation:
In Static Websites, Web pages are returned by the server which are prebuilt. They use simple languages such as HTML, CSS, or JavaScript.
There is no processing of content on the server (according to the user) in Static Websites. Web pages are returned by the server with no change therefore, static Websites are fast.
There is no interaction with databases.
Also, they are less costly as the host does not need to support server-side processing with different languages.
============
In Dynamic Websites, Web pages are returned by the server which are processed during runtime means they are not prebuilt web pages but they are built during runtime according to the user's demand.
These use server-side scripting languages such as PHP, Node.js, ASP.NET and many more supported by the server.
So, they are slower than static websites but updates and interaction with databases are possible.

NEW QUESTION 30
- (Exam Topic 1)
A company is launching a new application and will display application metrics on an Amazon CloudWatch dashboard. The company's product manager needs to access this dashboard periodically. The product manager does not have an AWS account. A solution architect must provide access to the product manager by following the principle of least privilege.
Which solution will meet these requirements?

A. Share the dashboard from the CloudWatch consol
B. Enter the product manager's email address, and complete the sharing step
C. Provide a shareable link for the dashboard to the product manager.
D. Create an IAM user specifically for the product manage
E. Attach the CloudWatch Read Only Access managed policy to the use
F. Share the new login credential with the product manage
G. Share the browser URL of the correct dashboard with the product manager.
H. Create an IAM user for the company's employees, Attach the View Only Access AWS managed policy to the IAM use
I. Share the new login credentials with the product manage
J. Ask the product manager to navigate to the CloudWatch console and locate the dashboard by name in the Dashboards section.
K. Deploy a bastion server in a public subne
L. When the product manager requires access to the dashboard, start the server and share the RDP credential
M. On the bastion server, ensure that the browser is configured to open the dashboard URL with cached AWS credentials that have appropriate permissions to view the dashboard.

Answer: B

NEW QUESTION 32
- (Exam Topic 1)
A company has a large Microsoft SharePoint deployment running on-premises that requires Microsoft Windows shared file storage. The company wants to migrate this workload to the AWS Cloud and is considering various storage options. The storage solution must be highly available and integrated with Active Directory for access control.
Which solution will satisfy these requirements?

A. Configure Amazon EFS storage and set the Active Directory domain for authentication
B. Create an SMB Me share on an AWS Storage Gateway tile gateway in two Availability Zones
C. Create an Amazon S3 bucket and configure Microsoft Windows Server to mount it as a volume
D. Create an Amazon FSx for Windows File Server file system on AWS and set the Active Directory domain for authentication

Answer: D

NEW QUESTION 37
- (Exam Topic 1)
A company hosts its web applications in the AWS Cloud. The company configures Elastic Load Balancers to use certificate that are imported into AWS Certificate Manager (ACM). The company's security team must be notified 30 days before the expiration of each certificate.
What should a solutions architect recommend to meet the requirement?

A. Add a rule m ACM to publish a custom message to an Amazon Simple Notification Service (Amazon SNS) topic every day beginning 30 days before any certificate will expire.
B. Create an AWS Config rule that checks for certificates that will expire within 30 day
C. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke a custom alert by way of Amazon Simple Notification Service (Amazon SNS) when AWS Config reports a noncompliant resource
D. Use AWS trusted Advisor to check for certificates that will expire within to day
E. Create an Amazon CloudWatch alarm that is based on Trusted Advisor metrics for check status changes Configure the alarm to send a custom alert by way of Amazon Simple rectification Service (Amazon SNS)
F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect any certificates that will expire within 30 day
G. Configure the rule to invoke an AWS Lambda functio
H. Configure the Lambda function to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS).

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/acm-certificate-expiration/

NEW QUESTION 38
- (Exam Topic 1)
A company hosts an application on multiple Amazon EC2 instances The application processes messages from an Amazon SQS queue writes to an Amazon RDS table and deletes the message from the queue Occasional duplicate records are found in the RDS table. The SQS queue does not contain any duplicate messages.
What should a solutions architect do to ensure messages are being processed once only?

A. Use the CreateQueue API call to create a new queue
B. Use the Add Permission API call to add appropriate permissions
C. Use the ReceiveMessage API call to set an appropriate wail time
D. Use the ChangeMessageVisibility APi call to increase the visibility timeout

Answer: D

Explanation:
The visibility timeout begins when Amazon SQS returns a message. During this time, the consumer processes and deletes the message. However, if the consumer fails before deleting the message and your system doesn't call the DeleteMessage action for that message before the visibility timeout expires, the message becomes visible to other consumers and the message is received again. If a message must be received only once, your consumer should delete it within the duration of the visibility timeout. https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html
Keyword: SQS queue writes to an Amazon RDS From this, Option D best suite & other Options ruled out [Option A - You can't intruduce one more Queue in the existing one; Option B - only Permission & Option C - Only Retrieves Messages] FIF O queues are designed to never introduce duplicate messages. However, your message producer might introduce duplicates in certain scenarios: for example, if the producer sends a message, does not receive a response, and then resends the same message. Amazon SQS APIs provide deduplication functionality that prevents your message producer from sending duplicates. Any duplicates introduced by the message producer are removed within a 5-minute deduplication interval. For standard queues, you might occasionally receive a duplicate copy of a message (at-least- once delivery). If you use a standard queue, you must design your applications to be idempotent (that is, they must not be affected adversely when processing the same message more than once).

NEW QUESTION 42
- (Exam Topic 1)
A company provides a Voice over Internet Protocol (VoIP) service that uses UDP connections. The service consists of Amazon EC2 instances that run in an Auto Scaling group. The company has deployments across multiple AWS Regions.
The company needs to route users to the Region with the lowest latency. The company also needs automated failover between Regions.
Which solution will meet these requirements?

A. Deploy a Network Load Balancer (NLB) and an associated target grou
B. Associate the target group with the Auto Scaling grou
C. Use the NLB as an AWS Global Accelerator endpoint in each Region.
D. Deploy an Application Load Balancer (ALB) and an associated target grou
E. Associate the target group with the Auto Scaling grou
F. Use the ALB as an AWS Global Accelerator endpoint in each Region.
G. Deploy a Network Load Balancer (NLB) and an associated target grou
H. Associate the target group with the Auto Scaling grou
I. Create an Amazon Route 53 latency record that points to aliases for each NL
J. Create an Amazon CloudFront distribution that uses the latency record as an origin.
K. Deploy an Application Load Balancer (ALB) and an associated target grou
L. Associate the target group with the Auto Scaling grou
M. Create an Amazon Route 53 weighted record that points to aliases for each AL
N. Deploy an Amazon CloudFront distribution that uses the weighted record as an origin.

Answer: D

Explanation:
https://aws.amazon.com/global-accelerator/faqs/
HTTP /HTTPS - ALB ; TCP and UDP - NLB; Lowest latency routing and more throughput. Also supports failover, uses Anycast Ip addressing - Global Accelerator Caching at Egde Locations – Cloutfront

WS Global Accelerator automatically checks the health of your applications and routes user traffic only to healthy application endpoints. If the health status changes or you make configuration updates, AWS Global Accelerator reacts instantaneously to route your users to the next available endpoint..

NEW QUESTION 43
- (Exam Topic 1)
A company runs its Infrastructure on AWS and has a registered base of 700.000 users for res document management application The company intends to create a product that converts large pdf files to jpg Imago files. The .pdf files average 5 MB in size. The company needs to store the original files and the converted files. A solutions architect must design a scalable solution to accommodate demand that will grow rapidly over lime.
Which solution meets these requirements MOST cost-effectively?

A. Save the pdf files to Amazon S3 Configure an S3 PUT event to invoke an AWS Lambda function to convert the files to jpg format and store them back in Amazon S3
B. Save the pdf files to Amazon DynamoD
C. Use the DynamoDB Streams feature to invoke an AWS Lambda function to convert the files to jpg format and store them hack in DynamoDB
D. Upload the pdf files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances.Amazon Elastic Block Store (Amazon EBS) storage and an Auto Scaling grou
E. Use a program In the EC2 instances to convert the files to jpg format Save the .pdf files and the .jpg files In the EBS store.
F. Upload the .pdf files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances, Amazon Elastic File System (Amazon EPS) storage, and an Auto Scaling grou
G. Use a program in the EC2 instances to convert the file to jpg format Save the pdf files and the jpg files in the EBS store.

Answer: A

Explanation:
Elastic BeanStalk is expensive, and DocumentDB has a 400KB max to upload files. So Lambda and S3 should be the one.

NEW QUESTION 48
- (Exam Topic 1)
A company hosts an application on AWS Lambda functions mat are invoked by an Amazon API Gateway API The Lambda functions save customer data to an Amazon Aurora MySQL database Whenever the company upgrades the database, the Lambda functions fail to establish database connections until the upgrade is complete The result is that customer data Is not recorded for some of the event
A solutions architect needs to design a solution that stores customer data that is created during database upgrades
Which solution will meet these requirements?

A. Provision an Amazon RDS proxy to sit between the Lambda functions and the database Configure the Lambda functions to connect to the RDS proxy
B. Increase the run time of me Lambda functions to the maximum Create a retry mechanism in the code that stores the customer data in the database
C. Persist the customer data to Lambda local storag
D. Configure new Lambda functions to scan the local storage to save the customer data to the database.
E. Store the customer data m an Amazon Simple Queue Service (Amazon SOS) FIFO queue Create a new Lambda function that polls the queue and stores the customer data in the database

Answer: D

Explanation:
https://www.learnaws.org/2020/12/13/aws-rds-proxy-deep-dive/
RDS proxy can improve application availability in such a situation by waiting for the new database instance to be functional and maintaining any requests received from the application during this time. The end result is that the application is more resilient to issues with the underlying database.
This will enable solution to hold data till the time DB comes back to normal. RDS proxy is to optimally utilize the connection between Lambda and DB. Lambda can open multiple connection concurrently which can be taxing on DB compute resources, hence RDS proxy was introduced to manage and leverage these connections efficiently.

NEW QUESTION 53
- (Exam Topic 1)
A solutions architect is designing a two-tier web application The application consists of a public-facing web tier hosted on Amazon EC2 in public subnets The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet Security is a high priority for the company
How should security groups be configured in this situation? (Select TWO)

A. Configure the security group for the web tier to allow inbound traffic on port 443 from 0.0.0.0/0.
B. Configure the security group for the web tier to allow outbound traffic on port 443 from 0.0.0.0/0.
C. Configure the security group for the database tier to allow inbound traffic on port 1433 from the securitygroup for the web tier.
D. Configure the security group for the database tier to allow outbound traffic on ports 443 and 1433 to the security group for the web tier.
E. Configure the security group for the database tier to allow inbound traffic on ports 443 and 1433 from the security group for the web tier.

Answer: AC

Explanation:
"Security groups create an outbound rule for every inbound rule." Not completely right. Statefull does NOT mean that if you create an inbound (or outbound) rule, it will create an outbound (or inbound) rule. What it does mean is: suppose you create an inbound rule on port 443 for the X ip. When a request enters on port 443 from X ip, it will allow traffic out for that request in the port 443. However, if you look at the outbound rules, there will not be any outbound rule on port 443 unless explicitly create it. In ACLs, which are stateless, you would have to create an inbound rule to allow incoming requests and an outbound rule to allow your application responds to those incoming requests.
https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#SecurityGroupRules

NEW QUESTION 58
- (Exam Topic 1)
A company recently migrated to AWS and wants to implement a solution to protect the traffic that flows in and out of the production VPC. The company had an inspection server in its on-premises data center. The inspection server performed specific operations such as traffic flow inspection and traffic filtering. The company wants to have the same functionalities in the AWS Cloud.
Which solution will meet these requirements?

A. Use Amazon GuardDuty for traffic inspection and traffic filtering in the production VPC
B. Use Traffic Mirroring to mirror traffic from the production VPC for traffic inspection and filtering.
C. Use AWS Network Firewall to create the required rules for traffic inspection and traffic filtering for the production VPC.
D. Use AWS Firewall Manager to create the required rules for traffic inspection and traffic filtering for the production VPC.

Answer: C

Explanation:
AWS Network Firewall supports both inspection and filtering as required

NEW QUESTION 60
- (Exam Topic 1)
A solutions architect is developing a multiple-subnet VPC architecture. The solution will consist of six subnets in two Availability Zones. The subnets are defined as public, private and dedicated for databases. Only the Amazon EC2 instances running in the private subnets should be able to access a database.
Which solution meets these requirements?

A. Create a now route table that excludes the route to the public subnets' CIDR block
B. Associate the route table to the database subnets.
C. Create a security group that denies ingress from the security group used by instances in the public subnet
D. Attach the security group to an Amazon RDS DB instance.
E. Create a security group that allows ingress from the security group used by instances in the private subnet
F. Attach the security group to an Amazon RDS DB instance.
G. Create a new peering connection between the public subnets and the private subnet
H. Create a different peering connection between the private subnets and the database subnets.

Answer: C

Explanation:
Security groups are stateful. All inbound traffic is blocked by default. If you create an inbound rule allowing traffic in, that traffic is automatically allowed back out again. You cannot block specific IP address using Security groups (instead use Network Access Control Lists).
"You can specify allow rules, but not deny rules." "When you first create a security group, it has no inbound rules. Therefore, no inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group." Source: https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#VPCSecurityGroups

NEW QUESTION 62
- (Exam Topic 1)
A company is storing sensitive user information in an Amazon S3 bucket The company wants to provide secure access to this bucket from the application tier running on Ama2on EC2 instances inside a VPC.
Which combination of steps should a solutions architect take to accomplish this? (Select TWO.)

A. Configure a VPC gateway endpoint for Amazon S3 within the VPC
B. Create a bucket policy to make the objects to the S3 bucket public
C. Create a bucket policy that limits access to only the application tier running in the VPC
D. Create an IAM user with an S3 access policy and copy the IAM credentials to the EC2 instance
E. Create a NAT instance and have the EC2 instances use the NAT instance to access the S3 bucket

Answer: AC

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-private-connection-no-authentication/

NEW QUESTION 64
- (Exam Topic 1)
A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed.
The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues.
Which solution will meet these requirements?

A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
B. Create an Amazon S3 File Gateway to extend the company's storage spac
C. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
D. Create an Amazon FSx for Windows File Server file system to extend the company's storage space.
E. Install a utility on each user's computer to access Amazon S3. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Answer: A

NEW QUESTION 66
- (Exam Topic 1)
A development team runs monthly resource-intensive tests on its general purpose Amazon RDS for MySQL DB instance with Performance Insights enabled. The testing lasts for 48 hours once a month and is the only process that uses the database. The team wants to reduce the cost of running the tests without reducing the compute and memory attributes of the DB instance.
Which solution meets these requirements MOST cost-effectively?

A. Stop the DB instance when tests are complete
B. Restart the DB instance when required.
C. Use an Auto Scaling policy with the DB instance to automatically scale when tests are completed.
D. Create a snapshot when tests are complete

E. Terminate the DB instance and restore the snapshot when required.
F. Modify the DB instance to a low-capacity instance when tests are complete
G. Modify the DB instance again when required.

Answer: A

NEW QUESTION 70
- (Exam Topic 1)
A company runs an online marketplace web application on AWS. The application serves hundreds of thousands of users during peak hours. The company needs a scalable, near-real-time solution to share the details of millions of financial transactions with several other internal applications Transactions also need to be processed to remove sensitive data before being stored in a document database for low-latency retrieval.
What should a solutions architect recommend to meet these requirements?

A. Store the transactions data into Amazon DynamoDB Set up a rule in DynamoDB to remove sensitive data from every transaction upon write Use DynamoDB Streams to share the transactions data with other applications
B. Stream the transactions data into Amazon Kinesis Data Firehose to store data in Amazon DynamoDB and Amazon S3 Use AWS Lambda integration with Kinesis Data Firehose to remove sensitive dat
C. Other applications can consume the data stored in Amazon S3
D. Stream the transactions data into Amazon Kinesis Data Streams Use AWS Lambda integration to remove sensitive data from every transaction and then store the transactions data in Amazon DynamoDB Other applications can consume the transactions data off the Kinesis data stream.
E. Store the batched transactions data in Amazon S3 as file
F. Use AWS Lambda to process every file and remove sensitive data before updating the files in Amazon S3 The Lambda function then stores the data in Amazon DynamoDB Other applications can consume transaction files stored in Amazon S3.

Answer: C

Explanation:
The destination of your Kinesis Data Firehose delivery stream. Kinesis Data Firehose can send data records to various destinations, including Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, and any HTTP endpoint that is owned by you or any of your third-party service providers. The following are the supported destinations:
* Amazon OpenSearch Service
* Amazon S3
* Datadog
* Dynatrace
* Honeycomb
* HTTP Endpoint
* Logic Monitor
* MongoDB Cloud
* New Relic
* Splunk
* Sumo Logic https://docs.aws.amazon.com/firehose/latest/dev/create-name.html https://aws.amazon.com/kinesis/data-streams/
Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. KDS can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events.

NEW QUESTION 75
- (Exam Topic 1)
A company needs to configure a real-time data ingestion architecture for its application. The company needs an API, a process that transforms data as the data is streamed, and a storage solution for the data.
Which solution will meet these requirements with the LEAST operational overhead?

A. Deploy an Amazon EC2 instance to host an API that sends data to an Amazon Kinesis data stream.Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data sourc
B. Use AWS Lambda functions to transform the dat
C. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
D. Deploy an Amazon EC2 instance to host an API that sends data to AWS Glu
E. Stop source/destination checking on the EC2 instanc
F. Use AWS Glue to transform the data and to send the data to Amazon S3.
G. Configure an Amazon API Gateway API to send data to an Amazon Kinesis data strea
H. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data sourc
I. Use AWS Lambda functions to transform the dat
J. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
K. Configure an Amazon API Gateway API to send data to AWS Glu
L. Use AWS Lambda functions to transform the dat
M. Use AWS Glue to send the data to Amazon S3.

Answer: C

NEW QUESTION 78
- (Exam Topic 1)
A company is implementing a new business application. The application runs on two Amazon EC2 instances and uses an Amazon S3 bucket for document storage. A solutions architect needs to ensure that the EC2 instances can access the S3 bucket.
What should the solutions architect do to meet this requirement?

A. Create an IAM role that grants access to the S3 bucke
B. Attach the role to the EC2 instances.
C. Create an IAM policy that grants access to the S3 bucke
D. Attach the policy to the EC2 instances.
E. Create an IAM group that grants access to the S3 bucke
F. Attach the group to the EC2 instances.

G. Create an IAM user that grants access to the S3 bucke
H. Attach the user account to the EC2 instances.

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/ec2-instance-access-s3-bucket/

NEW QUESTION 81
- (Exam Topic 1)
A company is preparing to deploy a new serverless workload. A solutions architect must use the principle of least privilege to configure permissions that will be used to run an AWS Lambda function. An Amazon EventBridge (Amazon CloudWatch Events) rule will invoke the function.
Which solution meets these requirements?

A. Add an execution role to the function with lambda:InvokeFunction as the action and * as the principal.
B. Add an execution role to the function with lambda:InvokeFunction as the action and Service:amazonaws.com as the principal.
C. Add a resource-based policy to the function with lambda:'* as the action and Service:events.amazonaws.com as the principal.
D. Add a resource-based policy to the function with lambda:InvokeFunction as the action and Service:events.amazonaws.com as the principal.

Answer: D

Explanation:
https://docs.aws.amazon.com/eventbridge/latest/userguide/resource-based-policies-eventbridge.html#lambda-pe

NEW QUESTION 86
- (Exam Topic 1)
A solutions architect is designing a VPC with public and private subnets. The VPC and subnets use IPv4 CIDR blocks. There is one public subnet and one private subnet in each of three Availability Zones (AZs) for high availability. An internet gateway is used to provide internet access for the public subnets. The private subnets require access to the internet to allow Amazon EC2 instances to download software updates.
What should the solutions architect do to enable Internet access for the private subnets?

A. Create three NAT gateways, one for each public subnet in each A
B. Create a private route table for each AZ that forwards non-VPC traffic to the NAT gateway in its AZ.
C. Create three NAT instances, one for each private subnet in each A
D. Create a private route table for each AZ that forwards non-VPC traffic to the NAT instance in its AZ.
E. Create a second internet gateway on one of the private subnet
F. Update the route table for the private subnets that forward non-VPC traffic to the private internet gateway.
G. Create an egress-only internet gateway on one of the public subnet
H. Update the route table for the private subnets that forward non-VPC traffic to the egress- only internet gateway.

Answer: A

Explanation:
https://aws.amazon.com/about-aws/whats-new/2018/03/introducing-amazon-vpc-nat-gateway-in-the-aws-govclo
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-comparison.html

NEW QUESTION 90
- (Exam Topic 1)
A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size.
Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again. The company also wants to automate remediation.
What should a solutions architect do to meet these requirements with the LEAST development effort?

A. Use an Amazon S3 bucket as a secure transfer poin
B. Use Amazon Inspector to scan me objects in the bucke
C. If objects contain Pl
D. trigger an S3 Lifecycle policy to remove the objects that contain Pll.
E. Use an Amazon S3 bucket as a secure transfer poin
F. Use Amazon Macie to scan the objects in the bucke
G. If objects contain Pl
H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects mat contain Pll.
I. Implement custom scanning algorithms in an AWS Lambda functio
J. Trigger the function when objects are loaded into the bucke
K. It objects contain RI
L. use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain Pll.
M. Implement custom scanning algorithms in an AWS Lambda functio
N. Trigger the function when objects are loaded into the bucke
O. If objects contain Pl
P. use Amazon Simple Email Service (Amazon STS) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects mot contain PII.

Answer: B

NEW QUESTION 91
- (Exam Topic 1)
A company has an Amazon S3 bucket that contains critical data. The company must protect the data from accidental deletion.
Which combination of steps should a solutions architect take to meet these requirements? (Choose two.)

A. Enable versioning on the S3 bucket.
B. Enable MFA Delete on the S3 bucket.
C. Create a bucket policy on the S3 bucket.
D. Enable default encryption on the S3 bucket.
E. Create a lifecycle policy for the objects in the S3 bucket.

Answer: AB

NEW QUESTION 96
- (Exam Topic 1)
A social media company allows users to upload images to its website. The website runs on Amazon EC2 instances. During upload requests, the website resizes the images to a standard size and stores the resized images in Amazon S3. Users are experiencing slow upload requests to the website.
The company needs to reduce coupling within the application and improve website performance. A solutions architect must design the most operationally efficient process for image uploads.
Which combination of actions should the solutions architect take to meet these requirements? (Choose two.)

A. Configure the application to upload images to S3 Glacier.
B. Configure the web server to upload the original images to Amazon S3.
C. Configure the application to upload images directly from each user's browser to Amazon S3 through the use of a presigned URL.
D. Configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploade
E. Use the function to resize the image
F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function on a schedule to resize uploaded images.

Answer: BD

NEW QUESTION 101
- (Exam Topic 1)
A company has applications that run on Amazon EC2 instances in a VPC. One of the applications needs to call the Amazon S3 API to store and read objects.
According to the company's security regulations, no traffic from the applications is allowed to travel across the internet.
Which solution will meet these requirements?

A. Configure an S3 interface endpoint.
B. Configure an S3 gateway endpoint.
C. Create an S3 bucket in a private subnet.
D. Create an S3 bucket in the same Region as the EC2 instance.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/privatelink-interface-endpoints.html#types-of-vpc-end
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html

NEW QUESTION 104
- (Exam Topic 1)
A company wants to move a multi-tiered application from on premises to the AWS Cloud to improve the application's performance. The application consists of application tiers that communicate with each other by way of RESTful services. Transactions are dropped when one tier becomes overloaded. A solutions architect must design a solution that resolves these issues and modernizes the application.
Which solution meets these requirements and is the MOST operationally efficient?

A. Use Amazon API Gateway and direct transactions to the AWS Lambda functions as the application laye
B. Use Amazon Simple Queue Service (Amazon SQS) as the communication layer between application services.
C. Use Amazon CloudWatch metrics to analyze the application performance history to determine the server's peak utilization during the performance failure
D. Increase the size of the application server's Amazon EC2 instances to meet the peak requirements.
E. Use Amazon Simple Notification Service (Amazon SNS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling grou
F. Use Amazon CloudWatch to monitor the SNS queue length and scale up and down as required.
G. Use Amazon Simple Queue Service (Amazon SQS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling grou
H. Use Amazon CloudWatch to monitor the SQS queue length and scale up when communication failures are detected.

Answer: A

Explanation:
https://aws.amazon.com/getting-started/hands-on/build-serverless-web-app-lambda-apigateway-s3-dynamodb-c Build a Serverless Web Application with AWS Lambda, Amazon API Gateway, AWS Amplify, Amazon DynamoDB, and Amazon Cognito. This example showed similar setup as question: Build a Serverless Web Application with AWS Lambda, Amazon API Gateway, AWS Amplify, Amazon DynamoDB, and Amazon Cognito

NEW QUESTION 108
- (Exam Topic 1)
A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The web application has static data and dynamic data. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53.
What should a solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins Configure Route 53 to route traffic to the CloudFront distribution.
B. Create an Amazon CloudFront distribution that has the ALB as an origin Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoin
C. Configure Route 53 to route traffic to the CloudFront distribution.
D. Create an Amazon CloudFront distribution that has the S3 bucket as an origin Create an AWS Global Accelerator standard accelerator that has the ALB and the CloudFront distribution as endpoints Create a custom domain name that points to the accelerator DNS name Use the custom domain name as an endpoint for the web application.

E. Create an Amazon CloudFront distribution that has the ALB as an origin
F. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint Create two domain name
G. Point one domain name to the CloudFront DNS name for dynamic content, Point the other domain name to the accelerator DNS name for static content Use the domain names as endpoints for the web application.

Answer: C

Explanation:
Static content can be cached at Cloud front Edge locations from S3 and dynamic content EC2 behind the ALB whose performance can be improved by Global Accelerator whose one endpoint is ALB and other Cloud front. So with regards to custom domain name endpoint is web application is R53 alias records for the custom domain point to web application
https://aws.amazon.com/blogs/networking-and-content-delivery/improving-availability-and-performance-for-ap

NEW QUESTION 109
- (Exam Topic 1)
A company uses Amazon S3 to store its confidential audit documents. The S3 bucket uses bucket policies to restrict access to audit team IAM user credentials according to the principle of least privilege. Company managers are worried about accidental deletion of documents in the S3 bucket and want a more secure solution.
What should a solutions architect do to secure the audit documents?

A. Enable the versioning and MFA Delete features on the S3 bucket.
B. Enable multi-factor authentication (MFA) on the IAM user credentials for each audit team IAM user account.
C. Add an S3 Lifecycle policy to the audit team's IAM user accounts to deny the s3:DeleteObject action during audit dates.
D. Use AWS Key Management Service (AWS KMS) to encrypt the S3 bucket and restrict audit team IAM user accounts from accessing the KMS key.

Answer: A

NEW QUESTION 113
- (Exam Topic 1)
A company has a data ingestion workflow that consists the following:

> An Amazon Simple Notification Service (Amazon SNS) topic for notifications about new data deliveries

> An AWS Lambda function to process the data and record metadata

The company observes that the ingestion workflow fails occasionally because of network connectivity issues. When such a failure occurs, the Lambda function does not ingest the corresponding data unless the company manually reruns the job.
Which combination of actions should a solutions architect take to ensure that the Lambda function ingests all data in the future? (Select TWO.)

A. Configure the Lambda function In multiple Availability Zones.
B. Create an Amazon Simple Queue Service (Amazon SQS) queue, and subscribe It to me SNS topic.
C. Increase the CPU and memory that are allocated to the Lambda function.
D. Increase provisioned throughput for the Lambda function.
E. Modify the Lambda function to read from an Amazon Simple Queue Service (Amazon SQS) queue

Answer: BE

NEW QUESTION 118
- (Exam Topic 1)
A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that the catalog is stored in a durable location.
What should a solutions architect do to meet these requirements?

A. Move the catalog to Amazon ElastiCache for Redis.
B. Deploy a larger EC2 instance with a larger instance store.
C. Move the catalog from the instance store to Amazon S3 Glacier Deep Archive.
D. Move the catalog to an Amazon Elastic File System (Amazon EFS) file system.

Answer: A

NEW QUESTION 122
- (Exam Topic 1)
A company stores call transcript files on a monthly basis. Users access the files randomly within 1 year of the call, but users access the files infrequently after 1 year. The company wants to optimize its solution by giving users the ability to query and retrieve files that are less than 1-year-old as quickly as possible. A delay in retrieving older files is acceptable.
Which solution will meet these requirements MOST cost-effectively?

A. Store individual files with tags in Amazon S3 Glacier Instant Retrieva
B. Query the tags to retrieve the files from S3 Glacier Instant Retrieval.
C. Store individual files in Amazon S3 Intelligent-Tierin
D. Use S3 Lifecycle policies to move the files to S3 Glacier Flexible Retrieval after 1 yea
E. Query and retrieve the files that are in Amazon S3 by using Amazon Athen
F. Query and retrieve the files that are in S3 Glacier by using S3 Glacier Select.
G. Store individual files with tags in Amazon S3 Standard storag
H. Store search metadata for each archive in Amazon S3 Standard storag
I. Use S3 Lifecycle policies to move the files to S3 Glacier Instant Retrieval after 1 yea
J. Query and retrieve the files by searching for metadata from Amazon S3.
K. Store individual files in Amazon S3 Standard storag
L. Use S3 Lifecycle policies to move the files to S3 Glacier Deep Archive after 1 yea
M. Store search metadata in Amazon RD
N. Query the files from Amazon RD

O. Retrieve the files from S3 Glacier Deep Archive.

Answer: B

Explanation:
"For archive data that needs immediate access, such as medical images, news media assets, or genomics data, choose the S3 Glacier Instant Retrieval storage class, an archive storage class that delivers the lowest cost storage with milliseconds retrieval. For archive data that does not require immediate access but needs the flexibility to retrieve large sets of data at no cost, such as backup or disaster recovery use cases, choose S3 Glacier Flexible Retrieval (formerly S3 Glacier), with retrieval in minutes or free bulk retrievals in 5-12 hours."
https://aws.amazon.com/about-aws/whats-new/2021/11/amazon-s3-glacier-instant-retrieval-storage-class/

NEW QUESTION 123
- (Exam Topic 1)
A company has an application that generates a large number of files, each approximately 5 MB in size. The files are stored in Amazon S3. Company policy requires the files to be stored for 4 years before they can be
deleted Immediate accessibility is always required as the files contain critical business data that is not easy to reproduce. The files are frequently accessed in the first 30 days of the object creation but are rarely accessed after the first 30 days
Which storage solution is MOST cost-effective?

A. Create an S3 bucket lifecycle policy to move Mm from S3 Standard to S3 Glacier 30 days from object creation Delete the Tiles 4 years after object creation
B. Create an S3 bucket lifecycle policy to move tiles from S3 Standard to S3 One Zone-infrequent Access (S3 One Zone-IA] 30 days from object creatio
C. Delete the fees 4 years after object creation
D. Create an S3 bucket lifecycle policy to move files from S3 Standard-infrequent Access (S3 Standard-IA) 30 from object creatio
E. Delete the ties 4 years after object creation
F. Create an S3 bucket Lifecycle policy to move files from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) 30 days from object creation Move the files to S3 Glacier 4 years after object carton.

Answer: B

Explanation:
https://aws.amazon.com/s3/storage-classes/?trk=66264cd8-3b73-416c-9693-ea7cf4fe846a&sc_channel=ps&s_k

NEW QUESTION 127
- (Exam Topic 1)
A company has a three-tier web application that is deployed on AWS. The web servers are deployed in a public subnet in a VPC. The application servers and database servers are deployed in private subnets in the same VPC. The company has deployed a third-party virtual firewall appliance from AWS Marketplace in an inspection VPC. The appliance is configured with an IP interface that can accept IP packets.
A solutions architect needs to Integrate the web application with the appliance to inspect all traffic to the application before the traffic teaches the web server.
Which solution will moot these requirements with the LEAST operational overhead?

A. Create a Network Load Balancer the public subnet of the application's VPC to route the traffic lo the appliance for packet inspection
B. Create an Application Load Balancer in the public subnet of the application's VPC to route the traffic to the appliance for packet inspection
C. Deploy a transit gateway m the inspection VPC Configure route tables to route the incoming pockets through the transit gateway
D. Deploy a Gateway Load Balancer in the inspection VPC Create a Gateway Load Balancer endpoint to receive the incoming packets and forward the packets to the appliance

Answer: D

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/scaling-network-traffic-inspection-using-aws-ga

NEW QUESTION 131
- (Exam Topic 1)
A company is deploying a new public web application to AWS. The application will run behind an Application Load Balancer (ALB). The application needs to be encrypted at the edge with an SSL/TLS certificate that is issued by an external certificate authority (CA). The certificate must be rotated each year before the certificate expires.
What should a solutions architect do to meet these requirements?

A. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificat
B. Apply the certificate to the AL
C. Use the managed renewal feature to automatically rotate the certificate.
D. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificat
E. Import the key material from the certificat
F. Apply the certificate to the AL
G. Use the managed renewal feature to automatically rotate the certificate.
H. Use AWS Certificate Manager (ACM) Private Certificate Authority to issue an SSL/TLS certificate from the root C
I. Apply the certificate to the AL
J. Use the managed renewal feature to automatically rotate the certificate.
K. Use AWS Certificate Manager (ACM) to import an SSL/TLS certificat
L. Apply the certificate to the AL
M. Use Amazon EventBridge (Amazon CloudWatch Events) to send a notification when the certificate is nearing expiratio
N. Rotate the certificate manually.

Answer: D

NEW QUESTION 134
- (Exam Topic 1)
A company wants to migrate its on-premises application to AWS. The application produces output files that vary in size from tens of gigabytes to hundreds of terabytes The application data must be stored in a standard file system structure The company wants a solution that scales automatically, is highly available, and

requires minimum operational overhead.
Which solution will meet these requirements?

A. Migrate the application to run as containers on Amazon Elastic Container Service (Amazon ECS) Use Amazon S3 for storage
B. Migrate the application to run as containers on Amazon Elastic Kubernetes Service (Amazon EKS) Use Amazon Elastic Block Store (Amazon EBS) for storage
C. Migrate the application to Amazon EC2 instances in a Multi-AZ Auto Scaling grou
D. Use Amazon Elastic File System (Amazon EFS) for storage.
E. Migrate the application to Amazon EC2 instances in a Multi-AZ Auto Scaling grou
F. Use Amazon Elastic Block Store (Amazon EBS) for storage.

Answer: C

Explanation:
EFS is a standard file system, it scales automatically and is highly available.

NEW QUESTION 137
- (Exam Topic 1)
A company has several web servers that need to frequently access a common Amazon RDS MySQL Multi-AZ DB instance The company wants a secure method for the web servers to connect to the database while meeting a security requirement to rotate user credentials frequently.
Which solution meets these requirements?

A. Store the database user credentials in AWS Secrets Manager Grant the necessary IAM permissions to allow the web servers to access AWS Secrets Manager
B. Store the database user credentials in AWS Systems Manager OpsCenter Grant the necessary IAM permissions to allow the web servers to access OpsCenter
C. Store the database user credentials in a secure Amazon S3 bucket Grant the necessary IAM permissions to allow the web servers to retrieve credentials and access the database.
D. Store the database user credentials in files encrypted with AWS Key Management Service (AWS KMS) on the web server file syste
E. The web server should be able to decrypt the files and access the database

Answer: A

Explanation:
AWS Secrets Manager helps you protect secrets needed to access your applications, services, and IT resources. The service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle.
https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html
Secrets Manager enables you to replace hardcoded credentials in your code, including passwords, with an API call to Secrets Manager to retrieve the secret programmatically. This helps ensure the secret can't be compromised by someone examining your code, because the secret no longer exists in the code. Also, you can configure Secrets Manager to automatically rotate the secret for you according to a specified schedule. This enables you to replace long-term secrets with short-term ones, significantly reducing the risk of compromise.

NEW QUESTION 139
- (Exam Topic 2)
A large media company hosts a web application on AWS. The company wants to start caching confidential media files so that users around the world will have reliable access to the files. The content is stored in Amazon S3 buckets. The company must deliver the content quickly, regardless of where the requests originate geographically.
Which solution will meet these requirements?

A. Use AWS DataSync to connect the S3 buckets to the web application.
B. Deploy AWS Global Accelerator to connect the S3 buckets to the web application.
C. Deploy Amazon CloudFront to connect the S3 buckets to CloudFront edge servers.
D. Use Amazon Simple Queue Service (Amazon SQS) to connect the S3 buckets to the web application.

Answer: C

Explanation:
CloudFront uses a local cache to provide the response, AWS Global accelerator proxies requests and connects to the application all the time for the response.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3

NEW QUESTION 143
- (Exam Topic 2)
A solutions architect needs to securely store a database user name and password that an application uses to access an Amazon RDS DB instance. The application that accesses the database runs on an Amazon EC2 instance. The solutions architect wants to create a secure parameter in AWS Systems Manager Parameter Store.
What should the solutions architect do to meet this requirement?

A. Create an IAM role that has read access to the Parameter Store paramete
B. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the paramete
C. Assign this IAM role to the EC2 instance.
D. Create an IAM policy that allows read access to the Parameter Store paramete
E. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the paramete
F. Assign this IAM policy to the EC2 instance.
G. Create an IAM trust relationship between the Parameter Store parameter and the EC2 instanc
H. Specify Amazon RDS as a principal in the trust policy.
I. Create an IAM trust relationship between the DB instance and the EC2 instanc
J. Specify Systems Manager as a principal in the trust policy.

Answer: B

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_aws-services-that-work-with-iam.html

NEW QUESTION 147
- (Exam Topic 2)
A company runs a global web application on Amazon EC2 instances behind an Application Load Balancer The application stores data in Amazon Aurora. The company needs to create a disaster recovery solution and can tolerate up to 30 minutes of downtime and potential data loss. The solution does not need to handle the load when the primary infrastructure is healthy
What should a solutions architect do to meet these requirements?

A. Deploy the application with the required infrastructure elements in place Use Amazon Route 53 to configure active-passive failover Create an Aurora Replica in a second AWS Region
B. Host a scaled-down deployment of the application in a second AWS Region Use Amazon Route 53 to configure active-active failover Create an Aurora Replica in the second Region
C. Replicate the primary infrastructure in a second AWS Region Use Amazon Route 53 to configure active-active failover Create an Aurora database that is restored from the latest snapshot
D. Back up data with AWS Backup Use the backup to create the required infrastructure in a second AWS Region Use Amazon Route 53 to configure active-passive failover Create an Aurora second primary instance in the second Region

Answer: A

Explanation:
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-types.html

NEW QUESTION 148
- (Exam Topic 2)
A company stores its application logs in an Amazon CloudWatch Logs log group. A new policy requires the company to store all application logs in Amazon OpenSearch Service (Amazon Elasticsearch Service) in near-real time.
Which solution will meet this requirement with the LEAST operational overhead?

A. Configure a CloudWatch Logs subscription to stream the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).
B. Create an AWS Lambda functio
C. Use the log group to invoke the function to write the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).
D. Create an Amazon Kinesis Data Firehose delivery strea
E. Configure the log group as the delivery stream's sourc
F. Configure Amazon OpenSearch Service (Amazon Elasticsearch Service) as the delivery stream's destination.
G. Install and configure Amazon Kinesis Agent on each application server to deliver the logs to Amazon Kinesis Data Stream
H. Configure Kinesis Data Streams to deliver the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service)

Answer: B

Explanation:
https://computingforgeeks.com/stream-logs-in-aws-from-cloudwatch-to-elasticsearch/

NEW QUESTION 149
- (Exam Topic 2)
A gaming company hosts a browser-based application on AWS. The users of the application consume a large number of videos and images that are stored in Amazon S3. This content is the same for all users.
The application has increased in popularity, and millions of users worldwide are accessing these media files. The company wants to provide the files to the users while reducing the load on the origin.
Which solution meets these requirements MOST cost-effectively?

A. Deploy an AWS Global Accelerator accelerator in front of the web servers.
B. Deploy an Amazon CloudFront web distribution in front of the S3 bucket.
C. Deploy an Amazon ElastiCache for Redis instance in front of the web servers.
D. Deploy an Amazon ElastiCache for Memcached instance in front of the web servers.

Answer: B

NEW QUESTION 152
- (Exam Topic 2)
A company is migrating its on-premises PostgreSQL database to Amazon Aurora PostgreSQL. The
on-premises database must remain online and accessible during the migration. The Aurora database must remain synchronized with the on-premises database.
Which combination of actions must a solutions architect take to meet these requirements? (Choose two.)

A. Create an ongoing replication task.
B. Create a database backup of the on-premises database
C. Create an AWS Database Migration Service (AWS DMS) replication server
D. Convert the database schema by using the AWS Schema Conversion Tool (AWS SCT).
E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor the database synchronization

Answer: CD

NEW QUESTION 155
- (Exam Topic 2)
An ecommerce company has an order-processing application that uses Amazon API Gateway and an AWS Lambda function. The application stores data in an Amazon Aurora PostgreSQL database. During a recent sales event, a sudden surge in customer orders occurred. Some customers experienced timeouts and the application did not process the orders of those customers A solutions architect determined that the CPU utilization and memory utilization were high on the database because of a large number of open connections The solutions architect needs to prevent the timeout errors while making the least possible changes to the application.
Which solution will meet these requirements?

A. Configure provisioned concurrency for the Lambda function Modify the database to be a global database in multiple AWS Regions
B. Use Amazon RDS Proxy to create a proxy for the database Modify the Lambda function to use the RDS Proxy endpoint instead of the database endpoint
C. Create a read replica for the database in a different AWS Region Use query string parameters in API Gateway to route traffic to the read replica
D. Migrate the data from Aurora PostgreSQL to Amazon DynamoDB by using AWS Database Migration Service (AWS DMS| Modify the Lambda function to use the OynamoDB table

Answer: D

NEW QUESTION 158
- (Exam Topic 2)
A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage.
The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand.
Which solution will meet these requirements?

A. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zone
B. Use an Amazon RDS DB instance in a Multi-AZ configuration.
C. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group in a single Availability Zon
D. Deploy the databaseon an EC2 instanc
E. Enable EC2 Auto Recovery.
F. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zone
G. Use an Amazon RDS DB instance with a read replica in a single Availability Zon
H. Promote the read replica to replace the primary DB instance if the primary DB instance fails.
I. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones Deploy the primary and secondary database servers on EC2 instances across multiple Availability Zones Use Amazon Elastic Block Store (Amazon EBS) Multi-Attach to create shared storage between the instances.

Answer: A

NEW QUESTION 163
- (Exam Topic 2)
A company is migrating an application from on-premises servers to Amazon EC2 instances. As part of the migration design requirements, a solutions architect must implement infrastructure metric alarms. The company does not need to take action if CPU utilization increases to more than 50% for a short burst of time.
However, if the CPU utilization increases to more than 50% and read IOPS on the disk are high at the same time, the company needs to act as soon as possible.
The solutions architect also must reduce false alarms.
What should the solutions architect do to meet these requirements?

A. Create Amazon CloudWatch composite alarms where possible.
B. Create Amazon CloudWatch dashboards to visualize the metrics and react to issues quickly.
C. Create Amazon CloudWatch Synthetics canaries to monitor the application and raise an alarm.
D. Create single Amazon CloudWatch metric alarms with multiple metric thresholds where possible.

Answer: A

NEW QUESTION 165
- (Exam Topic 2)
A company wants to direct its users to a backup static error page if the company's primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53. The domain is pointing to an Application Load Balancer (ALB). The company needs a solution that minimizes changes and infrastructure overhead.
Which solution will meet these requirements?

A. Update the Route 53 records to use a latency routing polic
B. Add a static error page that is hosted in an Amazon S3 bucket to the records so that the traffic is sent to the most responsive endpoints.
C. Set up a Route 53 active-passive failover configuratio
D. Direct traffic to a static error page that is hosted in an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.
E. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance that hosts a static error page as endpoint
F. Configure Route 53 to send requests to the instance only if the health checks fail for the ALB.
G. Update the Route 53 records to use a multivalue answer routing polic
H. Create a health chec
I. Direct traffic to the website if the health check passe
J. Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.

Answer: B

NEW QUESTION 168
- (Exam Topic 2)
A company has an AWS account used for software engineering. The AWS account has access to the company's on-premises data center through a pair of AWS Direct Connect connections. All non-VPC traffic routes to the virtual private gateway.
A development team recently created an AWS Lambda function through the console. The development team needs to allow the function to access a database that runs in a private subnet in the company's data center.
Which solution will meet these requirements?

A. Configure the Lambda function to run in the VPC with the appropriate security group.
B. Set up a VPN connection from AWS to the data cente
C. Route the traffic from the Lambda function through the VPN.
D. Update the route tables in the VPC to allow the Lambda function to access the on-premises data center through Direct Connect.
E. Create an Elastic IP addres
F. Configure the Lambda function to send traffic through the Elastic IP address without an elastic network interface.

Answer: A

Explanation:
https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html#vpc-managing-eni

NEW QUESTION 169
- (Exam Topic 2)
A solutions architect must design a solution that uses Amazon CloudFront with an Amazon S3 origin to store a static website. The company's security policy requires that all website traffic be inspected by AWS WAR
How should the solutions architect comply with these requirements?

A. Configure an S3 bucket policy lo accept requests coming from the AWS WAF Amazon Resource Name (ARN) only.
B. Configure Amazon CloudFront to forward all incoming requests to AWS WAF before requesting content from the S3 origin.
C. Configure a security group that allows Amazon CloudFront IP addresses to access Amazon S3 only.Associate AWS WAF to CloudFront.
D. Configure Amazon CloudFront and Amazon S3 to use an origin access identity (OAI) to restrict access to the S3 bucke
E. Enable AWS WAF on the distribution.

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-awswaf.html

NEW QUESTION 172
- (Exam Topic 2)
A company runs its ecommerce application on AWS. Every new order is published as a message in a RabbitMQ queue that runs on an Amazon EC2 instance in a single Availability Zone. These messages are processed by a different application that runs on a separate EC2 instance. This application stores the details in a PostgreSQL database on another EC2 instance. All the EC2 instances are in the same Availability Zone.
The company needs to redesign its architecture to provide the highest availability with the least operational overhead.
What should a solutions architect do to meet these requirements?

A. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
B. Create a Multi-AZ Auto Scaling group (or EC2 instances that host the applicatio
C. Create another Multi-AZAuto Scaling group for EC2 instances that host the PostgreSQL database.
D. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
E. Create a Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
F. Migrate the database to run ona Multi-AZ deployment of Amazon RDS for PostgreSQL.
G. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queu
H. Create another Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
I. Migrate the database to runon a Multi-AZ deployment of Amazon RDS fqjPostgreSQL.
J. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queu
K. Create another Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
L. Create a third Multi-AZ AutoScaling group for EC2 instances that host the PostgreSQL database.

Answer: B

NEW QUESTION 175
- (Exam Topic 2)
A company is planning to move its data to an Amazon S3 bucket. The data must be encrypted when it is stored in the S3 bucket. Additionally, the encryption key must be automatically rotated every year.
Which solution will meet these requirements with the LEAST operational overhead?

A. Move the data to the S3 bucke
B. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use the built-in key rotation behavior of SSE-S3 encryption keys.
C. Create an AWS Key Management Service {AWS KMS) customer managed ke
D. Enable automatic key rotatio
E. Set the S3 bucket's default encryption behavior to use the customer managed KMS ke
F. Move the data to the S3 bucket.
G. Create an AWS Key Management Service (AWS KMS) customer managed ke
H. Set the S3 bucket's default encryption behavior to use the customer managed KMS ke
I. Move the data to the S3 bucke
J. Manually rotate the KMS key every year.
K. Encrypt the data with customer key material before moving the data to the S3 bucke
L. Create an AWS Key Management Service (AWS KMS) key without key materia
M. Import the customer key material into the KMS ke
N. Enable automatic key rotation.

Answer: C

NEW QUESTION 176
- (Exam Topic 2)
A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should be protected throughout the entire application stack, and access to the information should be restricted to certain applications.
Which action should the solutions architect take?

A. Configure a CloudFront signed URL.
B. Configure a CloudFront signed cookie.
C. Configure a CloudFront field-level encryption profile.

D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/field-level-encryption.html
"With Amazon CloudFront, you can enforce secure end-to-end connections to origin servers by using HTTPS. Field-level encryption adds an additional layer of security that lets you protect specific data throughout system processing so that only certain applications can see it."

NEW QUESTION 179
- (Exam Topic 2)
A company wants to build a scalable key management Infrastructure to support developers who need to encrypt data in their applications.
What should a solutions architect do to reduce the operational burden?

A. Use multifactor authentication (MFA) to protect the encryption keys.
B. Use AWS Key Management Service (AWS KMS) to protect the encryption keys
C. Use AWS Certificate Manager (ACM) to create, store, and assign the encryption keys
D. Use an IAM policy to limit the scope of users who have access permissions to protect the encryption keys

Answer: B

Explanation:
https://aws.amazon.com/kms/faqs/#:~:text=If%20you%20are%20a%20developer%20who%20needs%20to%20d

NEW QUESTION 182
- (Exam Topic 2)
A company has two applications: a sender application that sends messages with payloads to be processed and a processing application intended to receive the messages with payloads. The company wants to implement an AWS service to handle messages between the two applications. The sender application can send about 1.000 messages each hour. The messages may take up to 2 days to be processed. If the messages fail to process, they must be retained so that they do not impact the processing of any remaining messages.
Which solution meets these requirements and is the MOST operationally efficient?

A. Set up an Amazon EC2 instance running a Redis databas
B. Configure both applications to use the instanc
C. Store, process, and delete the messages, respectively.
D. Use an Amazon Kinesis data stream to receive the messages from the sender applicatio
E. Integrate the processing application with the Kinesis Client Library (KCL).
F. Integrate the sender and processor applications with an Amazon Simple Queue Service (Amazon SQS) queu
G. Configure a dead-letter queue to collect the messages that failed to process.
H. Subscribe the processing application to an Amazon Simple Notification Service (Amazon SNS) topic to receive notifications to proces
I. Integrate the sender application to write to the SNS topic.

Answer: C

Explanation:
https://aws.amazon.com/blogs/compute/building-loosely-coupled-scalable-c-applications-with-amazon-sqs-and-
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-dead-letter-queues.htm

NEW QUESTION 187
- (Exam Topic 2)
A solutions architect needs to implement a solution to reduce a company's storage costs. All the company's data is in the Amazon S3 Standard storage class. The company must keep all data for at least 25 years. Data from the most recent 2 years must be highly available and immediately retrievable.
Which solution will meet these requirements?

A. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive immediately.
B. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 2 years.
C. Use S3 Intelligent-Tierin
D. Activate the archiving option to ensure that data is archived in S3 Glacier Deep Archive.
E. Set up an S3 Lifecycle policy to transition objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) immediately and to S3 Glacier Deep Archive after 2 years.

Answer: B

NEW QUESTION 190
- (Exam Topic 2)
A solutions architect needs to help a company optimize the cost of running an application on AWS. The application will use Amazon EC2 instances, AWS Fargate, and AWS Lambda for compute within the architecture.
The EC2 instances will run the data ingestion layer of the application. EC2 usage will be sporadic and unpredictable. Workloads that run on EC2 instances can be interrupted at any time. The application front end will run on Fargate, and Lambda will serve the API layer. The front-end utilization and API layer utilization will be predictable over the course of the next year.
Which combination of purchasing options will provide the MOST cost-effective solution for hosting this application? (Choose two.)

A. Use Spot Instances for the data ingestion layer
B. Use On-Demand Instances for the data ingestion layer
C. Purchase a 1-year Compute Savings Plan for the front end and API layer.
D. Purchase 1-year All Upfront Reserved instances for the data ingestion layer.
E. Purchase a 1-year EC2 instance Savings Plan for the front end and API layer.

Answer: AC

NEW QUESTION 193
- (Exam Topic 2)
A company is developing a file-sharing application that will use an Amazon S3 bucket for storage. The company wants to serve all the files through an Amazon CloudFront distribution. The company does not want the files to be accessible through direct navigation to the S3 URL.
What should a solutions architect do to meet these requirements?

A. Write individual policies for each S3 bucket to grant read permission for only CloudFront access.
B. Create an IAM use
C. Grant the user read permission to objects in the S3 bucke
D. Assign the user to CloudFront.
E. Write an S3 bucket policy that assigns the CloudFront distribution ID as the Principal and assigns the target S3 bucket as the Amazon Resource Name (ARN).
F. Create an origin access identity (OAI). Assign the OAI to the CloudFront distributio
G. Configure the S3 bucket permissions so that only the OAI has read permission.

Answer: D

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-access-to-amazon-s3/
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3

NEW QUESTION 198
- (Exam Topic 2)
An ecommerce company hosts its analytics application in the AWS Cloud. The application generates about 300 MB of data each month. The data is stored in JSON format. The company is evaluating a disaster recovery solution to back up the data. The data must be accessible in milliseconds if it is needed, and the data must be kept for 30 days.
Which solution meets these requirements MOST cost-effectively?

A. Amazon OpenSearch Service (Amazon Elasticsearch Service)
B. Amazon S3 Glacier
C. Amazon S3 Standard
D. Amazon RDS for PostgreSQL

Answer: C

NEW QUESTION 200
- (Exam Topic 2)
A company is planning to build a high performance computing (HPC) workload as a service solution that Is hosted on AWS A group of 16 AmazonEC2Ltnux
Instances requires the lowest possible latency for
node-to-node communication. The instances also need a shared block device volume for high-performing
storage.
Which solution will meet these requirements?

A. Use a duster placement grou
B. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon E BS) volume to all the instances by using Amazon EBS Multi-Attach
C. Use a cluster placement grou
D. Create shared 'lie systems across the instances by using Amazon Elastic File System (Amazon EFS)
E. Use a partition placement grou
F. Create shared tile systems across the instances by using Amazon Elastic File System (Amazon EFS).
G. Use a spread placement grou
H. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach

Answer: A

NEW QUESTION 204
- (Exam Topic 2)
A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-
northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet.
Which solutions will meet these requirements? (Choose two.)

A. Use AWS Control Tower to implement data residency guardrails to deny internet access and deny access to all AWS Regions except ap-northeast-3.
B. Use rules in AWS WAF to prevent internet acces
C. Deny access to all AWS Regions except ap-northeast-3 in the AWS account settings.
D. Use AWS Organizations to configure service control policies (SCPS) that prevent VPCs from gaining internet acces
E. Deny access to all AWS Regions except ap-northeast-3.
F. Create an outbound rule for the network ACL in each VPC to deny all traffic from 0.0.0.0/0. Create an IAM policy for each user to prevent the use of any AWS
Region other than ap-northeast-3.
G. Use AWS Config to activate managed rules to detect and alert for internet gateways and to detect and alert for new resources deployed outside of ap-
northeast-3.

Answer: AC

NEW QUESTION 205
- (Exam Topic 2)
A company hosts a website analytics application on a single Amazon EC2 On-Demand Instance. The analytics software is written in PHP and uses a MySQL
database. The analytics software, the web server that provides PHP, and the database server are all hosted on the EC2 instance. The application is showing signs
of performance degradation during busy times and is presenting 5xx errors. The company needs to make the application scale seamlessly.
Which solution will meet these requirements MOST cost-effectively?

A. Migrate the database to an Amazon RDS for MySQL DB instanc
B. Create an AMI of the web applicatio
C. Use the AMI to launch a second EC2 On-Demand Instanc
D. Use an Application Load Balancer to distribute the load to each EC2 instance.
E. Migrate the database to an Amazon RDS for MySQL DB instanc
F. Create an AMI of the web applicatio
G. Use the AMI to launch a second EC2 On-Demand Instanc
H. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.
I. Migrate the database to an Amazon Aurora MySQL DB instanc
J. Create an AWS Lambda function to stop the EC2 instance and change the instance typ
K. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization surpasses 75%.
L. Migrate the database to an Amazon Aurora MySQL DB instanc
M. Create an AMI of the web application.Apply the AMI to a launch templat
N. Create an Auto Scaling group with the launch template Configure the launch template to use a Spot Flee
O. Attach an Application Load Balancer to the Auto Scaling group.

Answer: D

NEW QUESTION 210
- (Exam Topic 2)
Organizers for a global event want to put daily reports online as static HTML pages. The pages are expected to generate millions of views from users around the world. The files are stored In an Amazon S3 bucket. A solutions architect has been asked to design an efficient and effective solution.
Which action should the solutions architect take to accomplish this?

A. Generate presigned URLs for the files.
B. Use cross-Region replication to all Regions.
C. Use the geoproximtty feature of Amazon Route 53.
D. Use Amazon CloudFront with the S3 bucket as its origin.

Answer: D

NEW QUESTION 211
- (Exam Topic 2)
A company's website provides users with downloadable historical performance reports. The website needs a solution that will scale to meet the company's website demands globally. The solution should be
cost-effective, limit the provisioning of infrastructure resources, and provide the fastest possible response time.
Which combination should a solutions architect recommend to meet these requirements?

A. Amazon CloudFront and Amazon S3
B. AWS Lambda and Amazon DynamoDB
C. Application Load Balancer with Amazon EC2 Auto Scaling
D. Amazon Route 53 with internal Application Load Balancers

Answer: A

Explanation:
Cloudfront for rapid response and s3 to minimize infrastructure.

NEW QUESTION 214
- (Exam Topic 2)
A company wants to migrate its existing on-premises monolithic application to AWS.
The company wants to keep as much of the front- end code and the backend code as possible. However, the company wants to break the application into smaller applications. A different team will manage each application. The company needs a highly scalable solution that minimizes operational overhead.
Which solution will meet these requirements?

A. Host the application on AWS Lambda Integrate the application with Amazon API Gateway.
B. Host the application with AWS Amplif
C. Connect the application to an Amazon API Gateway API that is integrated with AWS Lambda.
D. Host the application on Amazon EC2 instance
E. Set up an Application Load Balancer with EC2 instances in an Auto Scaling group as targets.
F. Host the application on Amazon Elastic Container Service (Amazon ECS) Set up an Application Load Balancer with Amazon ECS as the target.

Answer: D

Explanation:
https://aws.amazon.com/blogs/compute/microservice-delivery-with-amazon-ecs-and-application-load-balancers/

NEW QUESTION 217
- (Exam Topic 2)
A company has implemented a self-managed DNS solution on three Amazon EC2 instances behind a Network Load Balancer (NLB) in the us-west-2 Region. Most of the company's users are located in the United States and Europe. The company wants to improve the performance and availability of the solution. The company launches and configures three EC2 instances in the eu-west-1 Region and adds the EC2 instances as targets for a new NLB.
Which solution can the company use to route traffic to all the EC2 instances?

A. Create an Amazon Route 53 geolocation routing policy to route requests to one of the two NLB
B. Create an Amazon CloudFront distributio
C. Use the Route 53 record as the distribution's origin.
D. Create a standard accelerator in AWS Global Accelerato
E. Create endpoint groups in us-west-2 and eu-west-1. Add the two NLBs as endpoints for the endpoint groups.

F. Attach Elastic IP addresses to the six EC2 instance
G. Create an Amazon Route 53 geolocation routing policy to route requests to one of the six EC2 instance
H. Create an Amazon CloudFront distributio
I. Usethe Route 53 record as the distribution's origin.
J. Replace the two NLBs with two Application Load Balancers (ALBs). Create an Amazon Route 53 latency routing policy to route requests to one of the two ALB
K. Create an Amazon CloudFront distributio
L. Use the Route 53 record as the distribution's origin.

Answer: B

Explanation:
For standard accelerators, Global Accelerator uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure, which increases the availability of your applications. Endpoints for standard accelerators can be Network Load Balancers, Application Load Balancers, Amazon EC2 instances, or Elastic IP addresses that are located in one AWS Region or multiple Regions.
https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global-accelerator.html

NEW QUESTION 219
- (Exam Topic 2)
A company has a dynamic web application hosted on two Amazon EC2 instances. The company has its own SSL certificate, which is on each instance to perform SSL termination.
There has been an increase in traffic recently, and the operations team determined that SSL encryption and decryption is causing the compute capacity of the web servers to reach their maximum limit.
What should a solutions architect do to increase the application's performance?

A. Create a new SSL certificate using AWS Certificate Manager (ACM) install the ACM certificate on each instance
B. Create an Amazon S3 bucket Migrate the SSL certificate to the S3 bucket Configure the EC2 instances to reference the bucket for SSL termination
C. Create another EC2 instance as a proxy server Migrate the SSL certificate to the new instance and configure it to direct connections to the existing EC2 instances
D. Import the SSL certificate into AWS Certificate Manager (ACM) Create an Application Load Balancer with an HTTPS listener that uses the SSL certificate from ACM

Answer: D

Explanation:
https://aws.amazon.com/certificate-manager/:
"With AWS Certificate Manager, you can quickly request a certificate, deploy it on ACM-integrated AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let AWS Certificate Manager handle certificate renewals. It also enables you to create private certificates for
your internal resources and manage the certificate lifecycle centrally."

NEW QUESTION 222
- (Exam Topic 2)
A security team wants to limit access to specific services or actions in all of the team's AWS accounts. All accounts belong to a large organization in AWS Organizations. The solution must be scalable and there must be a single point where permissions can be maintained.
What should a solutions architect do to accomplish this?

A. Create an ACL to provide access to the services or actions.
B. Create a security group to allow accounts and attach it to user groups.
C. Create cross-account roles in each account to deny access to the services or actions.
D. Create a service control policy in the root organizational unit to deny access to the services or actions.

Answer: D

Explanation:
Service control policies (SCPs) are one type of policy that you can use to manage your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines. See
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html.

NEW QUESTION 223
- (Exam Topic 2)
A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing and new data by using SL. The company stores the data in an Amazon S3 bucket. The data requires encryption and must be replicated to a different AWS Region.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new S3 bucke
B. Load the data into the new S3 bucke
C. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
D. Use server-side encryption with AWS KMS multi-Region kays (SSE-KMS). Use Amazon Athena to query the data.
E. Create a new S3 bucke
F. Load the data into the new S3 bucke
G. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
H. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.
I. Load the data into the existing S3 bucke
J. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
K. Use server-side encryption with Amazon S3managed encryption keys (SSE-S3). Use Amazon Athena to query the data.
L. Load the data into the existing S3 bucke
M. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
N. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon RDS to query the data.

Answer: A

NEW QUESTION 225
- (Exam Topic 2)
A company wants to measure the effectiveness of its recent marketing campaigns. The company performs batch processing on csv files of sales data and stores the results «i an Amazon S3 bucket once every hour. The S3 bi petabytes of objects. The company runs one-time queries in Amazon Athena to determine which products are most popular on a particular date for a particular region Queries sometimes fail or take longer than expected to finish.
Which actions should a solutions architect take to improve the query performance and reliability? (Select TWO.)

A. Reduce the S3 object sizes to less than 126 MB
B. Partition the data by date and region n Amazon S3
C. Store the files as large, single objects in Amazon S3.
D. Use Amazon Kinosis Data Analytics to run the Queries as pan of the batch processing operation
E. Use an AWS duo extract, transform, and load (ETL) process to convert the csv files into Apache Parquet format.

Answer: CE

NEW QUESTION 226
- (Exam Topic 2)
A company runs a high performance computing (HPC) workload on AWS. The workload required low-latency network performance and high network throughput with tightly coupled node-to-node communication. The Amazon EC2 instances are properly sized for compute and storage capacity, and are launched using default options.
What should a solutions architect propose to improve the performance of the workload?

A. Choose a cluster placement group while launching Amazon EC2 instances.
B. Choose dedicated instance tenancy while launching Amazon EC2 instances.
C. Choose an Elastic Inference accelerator while launching Amazon EC2 instances.
D. Choose the required capacity reservation while launching Amazon EC2 instances.

Answer: A

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-placementgroup.html "A cluster placement group is a logical grouping of instances within a single Availability Zone that benefit from low network latency, high network throughput"

NEW QUESTION 229
- (Exam Topic 2)
A gaming company is designing a highly available architecture. The application runs on a modified Linux kernel and supports only UDP-based traffic. The company needs the front-end tier to provide the best possible user experience. That tier must have low latency, route traffic to the nearest edge location, and provide static IP addresses for entry into the application endpoints.
What should a solutions architect do to meet these requirements?

A. Configure Amazon Route 53 to forward requests to an Application Load Balance
B. Use AWS Lambda for the application in AWS Application Auto Scaling.
C. Configure Amazon CloudFront to forward requests to a Network Load Balance
D. Use AWS Lambda for the application in an AWS Application Auto Scaling group.
E. Configure AWS Global Accelerator to forward requests to a Network Load Balance
F. Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.
G. Configure Amazon API Gateway to forward requests to an Application Load Balance
H. Use AmazonEC2 instances for the application in an EC2 Auto Scaling group.

Answer: C

NEW QUESTION 233
- (Exam Topic 2)
A company needs to retain application logs files for a critical application for 10 years. The application team regularly accesses logs from the past month for troubleshooting, but logs older than 1 month are rarely accessed. The application generates more than 10 TB of logs per month.
Which storage option meets these requirements MOST cost-effectively?

A. Store the logs in Amazon S3 Use AWS Backup lo move logs more than 1 month old to S3 Glacier Deep Archive
B. Store the logs in Amazon S3 Use S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive
C. Store the logs in Amazon CloudWatch Logs Use AWS Backup to move logs more then 1 month old toS3 Glacier Deep Archive
D. Store the logs in Amazon CloudWatch Logs Use Amazon S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive

Answer: B

Explanation:
You need S3 to be able to archive the logs after one month. Cannot do that with CloudWatch Logs.

NEW QUESTION 237
- (Exam Topic 2)
A company has an ecommerce checkout workflow that writes an order to a database and calls a service to process the payment. Users are experiencing timeouts during the checkout process. When users resubmit the checkout form, multiple unique orders are created for the same desired transaction.
How should a solutions architect refactor this workflow to prevent the creation of multiple orders?

A. Configure the web application to send an order message to Amazon Kinesis Data Firehos
B. Set the payment service to retrieve the message from Kinesis Data Firehose and process the order.
C. Create a rule in AWS CloudTrail to invoke an AWS Lambda function based on the logged application path request Use Lambda to query the database, call the payment service, and pass in the order information.
D. Store the order in the databas

E. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to poll Amazon SN
F. retrieve the message, and process the order.
G. Store the order in the databas
H. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queu
I. Set the payment service to retrieve the message and process the orde
J. Delete the message from the queue.

Answer: D

Explanation:
This approach ensures that the order creation and payment processing steps are separate and atomic. By sending the order information to an SQS FIFO queue, the payment service can process the order one at a time and in the order they were received. If the payment service is unable to process an order, it can be retried later, preventing the creation of multiple orders. The deletion of the message from the queue after it is processed will prevent the same message from being processed multiple times.

NEW QUESTION 239
- (Exam Topic 2)
A company recently started using Amazon Aurora as the data store for its global ecommerce application When large reports are run developers report that the ecommerce application is performing poorly After reviewing metrics in Amazon CloudWatch, a solutions architect finds that the ReadIOPS and CPUUtilization metrics are spiking when monthly reports run.
What is the MOST cost-effective solution?

A. Migrate the monthly reporting to Amazon Redshift.
B. Migrate the monthly reporting to an Aurora Replica
C. Migrate the Aurora database to a larger instance class
D. Increase the Provisioned IOPS on the Aurora instance

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html
#Aurora.Replication.Replicas Aurora Replicas have two main purposes. You can issue queries to them to scale the read operations for your application. You typically do so by connecting to the reader endpoint of the cluster. That way, Aurora can spread the load for read-only connections across as many Aurora Replicas as you have in the cluster. Aurora Replicas also help to increase availability. If the writer instance in a cluster becomes unavailable, Aurora automatically promotes one of the reader instances to take its place as the new writer.
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Overview.html

NEW QUESTION 243
- (Exam Topic 2)
A corporation has recruited a new cloud engineer who should not have access to the CompanyConfidential Amazon S3 bucket. The cloud engineer must have read and write permissions on an S3 bucket named AdminTools.
Which IAM policy will satisfy these criteria?

A. Text, letter Description automatically generated

```
"Version": "2012-10-17",
"Statement": [
    {
        "Effect": "Allow",
        "Action": "s3:ListBucket",
        "Resource": "arn:aws:s3:::AdminTools"
    },
    {
        "Effect": "Allow",
        "Action": [ "s3:GetObject", "s3:PutObject" ],
        "Resource": "arn:aws:s3:::AdminTools/*"
    },
    {
        "Effect": "Deny",
        "Action": "s3:*",
        "Resource": [
            "arn:aws:s3:::CompanyConfidential/*",
            "arn:aws:s3:::CompanyConfidential"
        ]
    }
]
}
```

B. Text Description automatically generated

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": [
                "arn:aws:s3:::AdminTools",
                "arn:aws:s3:::CompanyConfidential/*"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": "arn:aws:s3:::CompanyConfidential"
        }
    ]
}
```

C. Text, application Description automatically generated

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::CompanyConfidential"
            ]
        }
    ]
}
```

D. Text, application Description automatically generated

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential",
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::AdminTools/*"
            ]
        }
    ]
}
```

Answer: A

Explanation:
https://docs.amazonaws.cn/en_us/IAM/latest/UserGuide/reference_policies_examples_s3_rw-bucket.html

NEW QUESTION 246
- (Exam Topic 2)
A company wants to manage Amazon Machine Images (AMIs). The company currently copies AMIs to the same AWS Region where the AMIs were created. The company needs to design an application that captures AWS API calls and sends alerts whenever the Amazon EC2 CreateImage API operation is called within the company's account.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS Lambda function to query AWS CloudTrail logs and to send an alert when a CreateImage API call is detected.
B. Configure AWS CloudTrail with an Amazon Simple Notification Service {Amazon SNS) notification that occurs when updated logs are sent to Amazon S3. Use Amazon Athena to create a new table and to query on CreateImage when an API call is detected.

C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for the CreateImage API call.Configure the target as an Amazon Simple Notification Service (Amazon SNS) topic to send an alert when a CreateImage API call is detected.
D. Configure an Amazon Simple Queue Service (Amazon SQS) FIFO queue as a target for AWS CloudTrail log
E. Create an AWS Lambda function to send an alert to an Amazon Simple Notification Service (Amazon SNS) topic when a CreateImage API call is detected.

Answer: B

NEW QUESTION 248
- (Exam Topic 2)
A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost.
Which solution meets the availability requirement MOST cost-effectively?

A. Use all EC2 Spot Instance
B. Stop the RDS database when it is not in use.
C. Purchase EC2 Instance Savings Plans to cover five EC2 instance
D. Purchase an RDS Reserved DB Instance
E. Purchase two EC2 Reserved Instances Use up to three additional EC2 Spot Instances as neede
F. Stop the RDS database when it is not in use.
G. Purchase EC2 Instance Savings Plans to cover two EC2 instance
H. Use up to three additional EC2 On-Demand Instances as neede
I. Purchase an RDS Reserved DB Instance.

Answer: D

NEW QUESTION 252
- (Exam Topic 2)
A company sells ringtones created from clips of popular songs. The files containing the ringtones are stored in Amazon S3 Standard and are at least 128 KB in size. The company has millions of files, but downloads are infrequent for ringtones older than 90 days. The company needs to save money on storage while keeping the most accessed files readily available for its users.
Which action should the company take to meet these requirements MOST cost-effectively?

A. Configure S3 Standard-Infrequent Access (S3 Standard-IA) storage for the initial storage tier of the objects.
B. Move the files to S3 Intelligent-Tiering and configure it to move objects to a less expensive storage tier after 90 days.
C. Configure S3 inventory to manage objects and move them to S3 Standard-Infrequent Access (S3 Standard-1A) after 90 days.
D. Implement an S3 Lifecycle policy that moves the objects from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-1A) after 90 days.

Answer: D

NEW QUESTION 257
- (Exam Topic 3)
A company s order system sends requests from clients to Amazon EC2 instances The EC2 instances process the orders and men store the orders in a database on Amazon RDS Users report that they must reprocess orders when the system fails. The company wants a resilient solution that can process orders automatically it a system outage occurs.
What should a solutions architect do to meet these requirements?

A. Move (he EC2 Instances into an Auto Scaling group Create an Amazon EventBridge (Amazon CloudWatch Events) rule to target an Amazon Elastic Container Service (Amazon ECS) task
B. Move the EC2 instances into an Auto Scaling group behind an Application Load Balancer (ALB) Update the order system to send messages to the ALB endpoint.
C. Move the EC2 instances into an Auto Scaling group Configure the order system to send messages to an Amazon Simple Queue Service (Amazon SQS) queue Configure the EC2 instances to consume messages from the queue
D. Create an Amazon Simple Notification Service (Amazon SNS) topic Create an AWS Lambda function, and subscribe the function to the SNS topic Configure the order system to send messages to the SNS topic Send a command to the EC2 instances to process the messages by using AWS Systems Manager Run Command

Answer: C

NEW QUESTION 259
- (Exam Topic 3)
A company hosts a marketing website in an on-premises data center. The website consists of static documents and runs on a single server. An administrator updates the website content infrequently and uses an SFTP client to upload new documents.
The company decides to host its website on AWS and to use Amazon CloudFront. The company's solutions architect creates a CloudFront distribution. The solutions architect must design the most cost-effective and resilient architecture for website hosting to serve as the CloudFront origin.
Which solution will meet these requirements?

A. Create a virtual server by using Amazon Lightsai
B. Configure the web server in the Lightsail instance.Upload website content by using an SFTP client.
C. Create an AWS Auto Scaling group for Amazon EC2 instance
D. Use an Application Load Balancer.Upload website content by using an SFTP client.
E. Create a private Amazon S3 bucke
F. Use an S3 bucket policy to allow access from a CloudFront origin access identity (OAI). Upload website content by using theAWSCLI.
G. Create a public Amazon S3 bucke
H. Configure AWS Transfer for SFT
I. Configure the S3 bucket for website hostin
J. Upload website content by using the SFTP client.

Answer: C

Explanation:

https://docs.aws.amazon.com/cli/latest/reference/transfer/describe-server.html

NEW QUESTION 263

......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Solution-Architect-Associate Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Solution-Architect-Associate Product From:

https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/

Money Back Guarantee

AWS-Solution-Architect-Associate Practice Exam Features:

* AWS-Solution-Architect-Associate Questions and Answers Updated Frequently

* AWS-Solution-Architect-Associate Practice Questions Verified by Expert Senior Certified Staff

* AWS-Solution-Architect-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year