

Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



NEW QUESTION 1

- (Topic 1)

A company hosts a containerized web application on a fleet of on-premises servers that process incoming requests. The number of requests is growing quickly. The on-premises servers cannot handle the increased number of requests. The company wants to move the application to AWS with minimum code changes and minimum development effort.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scaling
- B. Use an Application Load Balancer to distribute the incoming requests.
- C. Use two Amazon EC2 instances to host the containerized web application
- D. Use an Application Load Balancer to distribute the incoming requests
- E. Use AWS Lambda with a new code that uses one of the supported languages
- F. Create multiple Lambda functions to support the load
- G. Use Amazon API Gateway as an entry point to the Lambda functions.
- H. Use a high performance computing (HPC) solution such as AWS ParallelCluster to establish an HPC cluster that can process the incoming requests at the appropriate scale.

Answer: A

Explanation:

AWS Fargate is a serverless compute engine that lets users run containers without having to manage servers or clusters of Amazon EC2 instances¹. Users can use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scaling. Amazon ECS is a fully managed container orchestration service that supports both Docker and Kubernetes². Service Auto Scaling is a feature that allows users to adjust the desired number of tasks in an ECS service based on CloudWatch metrics, such as CPU utilization or request count³. Users can use AWS Fargate on Amazon ECS to migrate the application to AWS with minimum code changes and minimum development effort, as they only need to package their application in containers and specify the CPU and memory requirements.

Users can also use an Application Load Balancer to distribute the incoming requests. An Application Load Balancer is a load balancer that operates at the application layer and routes traffic to targets based on the content of the request. Users can register their ECS tasks as targets for an Application Load Balancer and configure listener rules to route requests to different target groups based on path or host headers. Users can use an Application Load Balancer to improve the availability and performance of their web application.

NEW QUESTION 2

- (Topic 1)

A company hosts its web applications in the AWS Cloud. The company configures Elastic Load Balancers to use certificates 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 in 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 days
- 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 30 days
- 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 Notification Service (Amazon SNS)
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect any certificates that will expire within 30 days
- G. Configure the rule to invoke an AWS Lambda function
- 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 3

- (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 point
- B. Use Amazon Inspector to scan the objects in the bucket
- C. If objects contain PII
- D. Trigger an S3 Lifecycle policy to remove the objects that contain PII.
- E. Use an Amazon S3 bucket as a secure transfer point
- F. Use Amazon Macie to scan the objects in the bucket
- G. If objects contain PII
- H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- I. Implement custom scanning algorithms in an AWS Lambda function
- J. Trigger the function when objects are loaded into the bucket
- K. If objects contain PII
- L. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- M. Implement custom scanning algorithms in an AWS Lambda function
- N. Trigger the function when objects are loaded into the bucket
- O. If objects contain PII

P. use Amazon Simple Email Service (Amazon SES) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects not contain PII.

Answer: B

Explanation:

To meet the requirements of detecting and alerting the administrators when PII is shared and automating remediation with the least development effort, the best approach would be to use Amazon S3 bucket as a secure transfer point and scan the objects in the bucket with Amazon Macie. Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data stored in Amazon S3. It can be used to classify sensitive data, monitor access to sensitive data, and automate remediation actions.

In this scenario, after uploading the files to the Amazon S3 bucket, the objects can be scanned for PII by Amazon Macie, and if it detects any PII, it can trigger an Amazon Simple Notification Service (SNS) notification to alert the administrators to remove the objects containing PII. This approach requires the least development effort, as Amazon Macie already has pre-built data classification rules that can detect PII in various formats. Hence, option B is the correct answer.

References:

? Amazon Macie User Guide: <https://docs.aws.amazon.com/macie/latest/userguide/what-is-macie.html>

? AWS Well-Architected Framework - Security Pillar: <https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/welcome.html>

NEW QUESTION 4

- (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 reaches the web server. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a Network Load Balancer in the public subnet of the application's VPC to route the traffic to 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 in the inspection VPC. Configure route tables to route the incoming packets 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-gateway-load-balancer/>

NEW QUESTION 5

- (Topic 1)

A company runs an on-premises application that is powered by a MySQL database. The company is migrating the application to AWS to increase the application's elasticity and availability.

The current architecture shows heavy read activity on the database during times of normal operation. Every 4 hours, the company's development team pulls a full export of the production database to populate a database in the staging environment. During this period, users experience unacceptable application latency. The development team is unable to use the staging environment until the procedure completes.

A solutions architect must recommend a replacement architecture that alleviates the application latency issue. The replacement architecture also must give the development team the ability to continue using the staging environment without delay.

Which solution meets these requirements?

- A. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production.
- B. Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.
- C. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production. Use database cloning to create the staging database on-demand.
- D. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production. Use the standby instance for the staging database.
- E. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production.
- F. Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/aws/amazon-aurora-fast-database-cloning/>

NEW QUESTION 6

- (Topic 1)

A company wants to run its critical applications in containers to meet requirements for scalability and availability. The company prefers to focus on maintenance of the critical applications. The company does not want to be responsible for provisioning and managing the underlying infrastructure that runs the containerized workload.

What should a solutions architect do to meet those requirements?

- A. Use Amazon EC2 Instances, and install Docker on the instances.
- B. Use Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 worker nodes.
- C. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate.
- D. Use Amazon EC2 instances from an Amazon Elastic Container Service (Amazon ECS)-optimized Amazon Machine Image (AMI).

Answer: C

Explanation:

using AWS ECS on AWS Fargate since the requirements are for scalability and availability without having to provision and manage the underlying infrastructure to run the containerized workload. <https://docs.aws.amazon.com/AmazonECS/latest/userguide/what-is-fargate.html>

NEW QUESTION 7

- (Topic 1)

A company hosts an application on AWS Lambda functions that 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 the 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 storage.
- D. Configure new Lambda functions to scan the local storage to save the customer data to the database.
- E. Store the customer data in an Amazon Simple Queue Service (Amazon SQS) 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 8

- (Topic 1)

A company runs an ecommerce application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. The Auto Scaling group scales based on CPU utilization metrics. The ecommerce application stores the transaction data in a MySQL 8.0 database that is hosted on a large EC2 instance.

The database's performance degrades quickly as application load increases. The application handles more read requests than write transactions. The company wants a solution that will automatically scale the database to meet the demand of unpredictable read workloads while maintaining high availability.

Which solution will meet these requirements?

- A. Use Amazon Redshift with a single node for leader and compute functionality.
- B. Use Amazon RDS with a Single-AZ deployment. Configure Amazon RDS to add reader instances in a different Availability Zone.
- C. Use Amazon Aurora with a Multi-AZ deployment.
- D. Configure Aurora Auto Scaling with Aurora Replicas.
- E. Use Amazon ElastiCache for Memcached with EC2 Spot Instances.

Answer: C

Explanation:

AURORA is 5x performance improvement over MySQL on RDS and handles more read requests than write,; maintaining high availability = Multi-AZ deployment

NEW QUESTION 9

- (Topic 1)

A company has a website hosted on AWS. The website is behind an Application Load Balancer (ALB) that is configured to handle HTTP and HTTPS separately. The company wants to forward all requests to the website so that the requests will use HTTPS.

What should a solutions architect do to meet this requirement?

- A. Update the ALB's network ACL to accept only HTTPS traffic.
- B. Create a rule that replaces the HTTP in the URL with HTTPS.
- C. Create a listener rule on the ALB to redirect HTTP traffic to HTTPS.
- D. Replace the ALB with a Network Load Balancer configured to use Server Name Indication (SNI).

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https-using-alb/>

How can I redirect HTTP requests to HTTPS using an Application Load Balancer? Last updated: 2020-10-30 I want to redirect HTTP requests to HTTPS using Application Load Balancer listener rules. How can I do this? Resolution Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https-using-alb/>

NEW QUESTION 10

- (Topic 1)

An image-processing company has a web application that users use to upload images. The application uploads the images into an Amazon S3 bucket. The company has set up S3 event notifications to publish the object creation events to an Amazon Simple Queue Service (Amazon SQS) standard queue. The SQS queue serves as the event source for an AWS Lambda function that processes the images and sends the results to users through email.

Users report that they are receiving multiple email messages for every uploaded image. A solutions architect determines that SQS messages are invoking the Lambda function more than once, resulting in multiple email messages.

What should the solutions architect do to resolve this issue with the LEAST operational overhead?

- A. Set up long polling in the SQS queue by increasing the ReceiveMessage wait time to 30 seconds.
- B. Change the SQS standard queue to an SQS FIFO queue.
- C. Use the message deduplication ID to discard duplicate messages.
- D. Increase the visibility timeout in the SQS queue to a value that is greater than the total of the function timeout and the batch window timeout.
- E. Modify the Lambda function to delete each message from the SQS queue immediately after the message is read before processing.

Answer: C

NEW QUESTION 10

- (Topic 1)

An ecommerce company wants to launch a one-deal-a-day website on AWS. Each day will feature exactly one product on sale for a period of 24 hours. The company wants to be able to handle millions of requests each hour with millisecond latency during peak hours.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon S3 to host the full website in different S3 buckets Add Amazon CloudFront distributions Set the S3 buckets as origins for the distributions Store the order data in Amazon S3
- B. Deploy the full website on Amazon EC2 instances that run in Auto Scaling groups across multiple Availability Zones Add an Application Load Balancer (ALB) to distribute the website traffic Add another ALB for the backend APIs Store the data in Amazon RDS for MySQL
- C. Migrate the full application to run in containers Host the containers on Amazon Elastic Kubernetes Service (Amazon EKS) Use the Kubernetes Cluster Autoscaler to increase and decrease the number of pods to process bursts in traffic Store the data in Amazon RDS for MySQL
- D. Use an Amazon S3 bucket to host the website's static content Deploy an Amazon CloudFront distributio
- E. Set the S3 bucket as the origin Use Amazon API Gateway and AWS Lambda functions for the backend APIs Store the data in Amazon DynamoDB

Answer: D

Explanation:

To launch a one-deal-a-day website on AWS with millisecond latency during peak hours and with the least operational overhead, the best option is to use an Amazon S3 bucket to host the website's static content, deploy an Amazon CloudFront distribution, set the S3 bucket as the origin, use Amazon API Gateway and AWS Lambda functions for the backend APIs, and store the data in Amazon DynamoDB. This option requires minimal operational overhead and can handle millions of requests each hour with millisecond latency during peak hours. Therefore, option D is the correct answer.

Reference: <https://aws.amazon.com/blogs/compute/building-a-serverless-multi-player-game-with-aws-lambda-and-amazon-dynamodb/>

NEW QUESTION 13

- (Topic 1)

A company hosts a data lake on AWS. The data lake consists of data in Amazon S3 and Amazon RDS for PostgreSQL. The company needs a reporting solution that provides data visualization and includes all the data sources within the data lake. Only the company's management team should have full access to all the visualizations. The rest of the company should have only limited access.

Which solution will meet these requirements?

- A. Create an analysis in Amazon QuickSight
- B. Connect all the data sources and create new dataset
- C. Publish dashboards to visualize the dat
- D. Share the dashboards with the appropriate IAM roles.
- E. Create an analysis in Amazon QuickSight
- F. Connect all the data sources and create new dataset
- G. Publish dashboards to visualize the dat
- H. Share the dashboards with the appropriate users and groups.
- I. Create an AWS Glue table and crawler for the data in Amazon S3. Create an AWS Glue extract, transform, and load (ETL) job to produce report
- J. Publish the reports to Amazon S3. Use S3 bucket policies to limit access to the reports.
- K. Create an AWS Glue table and crawler for the data in Amazon S3. Use Amazon Athena Federated Query to access data within Amazon RDS for PostgreSQL
- L. Generate reports by using Amazon Athena
- M. Publish the reports to Amazon S3. Use S3 bucket policies to limit access to the reports.

Answer: B

Explanation:

Amazon QuickSight is a data visualization service that allows you to create interactive dashboards and reports from various data sources, including Amazon S3 and Amazon RDS for PostgreSQL. You can connect all the data sources and create new datasets in QuickSight, and then publish dashboards to visualize the data. You can also share the dashboards with the appropriate users and groups, and control their access levels using IAM roles and permissions.

Reference: <https://docs.aws.amazon.com/quicksight/latest/user/working-with-data-sources.html>

NEW QUESTION 15

- (Topic 1)

A company is implementing a shared storage solution for a media application that is hosted in the AWS Cloud. The company needs the ability to use SMB clients to access data. The solution must be fully managed.

Which AWS solution meets these requirements?

- A. Create an AWS Storage Gateway volume gateway
- B. Create a file share that uses the required client protocol. Connect the application server to the file share.
- C. Create an AWS Storage Gateway tape gateway. Configure (apes to use Amazon S3. Connect the application server to the tape gateway
- D. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance.
- E. Connect the application server to the file share.
- F. Create an Amazon FSx for Windows File Server file system. Attach the file system to the origin server.
- G. Connect the application server to the file system.

Answer: D

Explanation:

<https://aws.amazon.com/fsx/lustre/>

Amazon FSx has native support for Windows file system features and for the industry-standard Server Message Block (SMB) protocol to access file storage over a network. <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>

NEW QUESTION 16

- (Topic 1)

A company runs multiple Windows workloads on AWS. The company's employees use Windows file shares that are hosted on two Amazon EC2 instances. The file shares synchronize data between themselves and maintain duplicate copies. The company wants a highly available and durable storage solution that preserves how users currently access the files.

What should a solutions architect do to meet these requirements?

- A. Migrate all the data to Amazon S3 Set up IAM authentication for users to access files
- B. Set up an Amazon S3 File Gateway
- C. Mount the S3 File Gateway on the existing EC2 Instances.
- D. Extend the file share environment to Amazon FSx for Windows File Server with a Multi- AZ configuratio
- E. Migrate all the data to FSx for Windows File Server.
- F. Extend the file share environment to Amazon Elastic File System (Amazon EFS) with a Multi-AZ configuratio
- G. Migrate all the data to Amazon EFS.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/AmazonEFS.html> Amazon FSx for Windows File Server provides fully managed Microsoft Windows file servers, backed by a fully native Windows file system. <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>

NEW QUESTION 21

- (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 wait 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] FIFO 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 22

- (Topic 1)

A company is planning to use an Amazon DynamoDB table for data storage. The company is concerned about cost optimization. The table will not be used on most mornings. In the evenings, the read and write traffic will often be unpredictable. When traffic spikes occur, they will happen very quickly.

What should a solutions architect recommend?

- A. Create a DynamoDB table in on-demand capacity mode.
- B. Create a DynamoDB table with a global secondary index.
- C. Create a DynamoDB table with provisioned capacity and auto scaling.
- D. Create a DynamoDB table in provisioned capacity mode, and configure it as a global table.

Answer: A

NEW QUESTION 24

- (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

Explanation:

https://www.amazonaws.cn/en/certificate-manager/faqs/#Managed_renewal_and_deployment

NEW QUESTION 29

- (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-permissions>

NEW QUESTION 33

- (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 new 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 35

- (Topic 1)

A company is using a SQL database to store movie data that is publicly accessible. The database runs on an Amazon RDS Single-AZ DB instance. A script runs queries at random intervals each day to record the number of new movies that have been added to the database. The script must report a final total during business hours. The company's development team notices that the database performance is inadequate for development tasks when the script is running. A solutions architect must recommend a solution to resolve this issue. Which solution will meet this requirement with the LEAST operational overhead?

- A. Modify the DB instance to be a Multi-AZ deployment
- B. Create a read replica of the database. Configure the script to query only the read replica.
- C. Instruct the development team to manually export the entries in the database at the end of each day.
- D. Use Amazon ElastiCache to cache the common queries that the script runs against the database.

Answer: B

NEW QUESTION 36

- (Topic 1)

A company is migrating applications to AWS. The applications are deployed in different accounts. The company manages the accounts centrally by using AWS Organizations. The company's security team needs a single sign-on (SSO) solution across all the company's accounts. The company must continue managing the users and groups in its on-premises self-managed Microsoft Active Directory.

Which solution will meet these requirements?

- A. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console.
- B. Create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- C. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console.
- D. Create a two-way forest trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- E. Use AWS Directory Service.
- F. Create a two-way trust relationship with the company's self-managed Microsoft Active Directory.
- G. Deploy an identity provider (IdP) on-premise.
- H. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console.

Answer: A

Explanation:

To provide single sign-on (SSO) across all the company's accounts while continuing to manage users and groups in its on-premises self-managed Microsoft Active Directory, the solution is to enable AWS Single Sign-On (SSO) from the AWS SSO console and create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory. This solution is described in the AWS documentation.

NEW QUESTION 41

- (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 auditteam 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 42

- (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-govcloud-us-region/#:~:text=NAT%20Gateway%20is%20a%20highly,instances%20in%20a%20private%20subnet.>

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-comparison.html>

NEW QUESTION 45

- (Topic 1)

A company has an application that runs on Amazon EC2 instances and uses an Amazon

Aurora database. The EC2 instances connect to the database by using user names and passwords that are stored locally in a file. The company wants to minimize the operational overhead of credential management.

What should a solutions architect do to accomplish this goal?

- A. Use AWS Secrets Manage
- B. Turn on automatic rotation.
- C. Use AWS Systems Manager Parameter Stor
- D. Turn on automatic rotation.
- E. Create an Amazon S3 bucket lo store objects that are encrypted with an AWS Key
- F. Management Service (AWS KMS) encryption ke
- G. Migrate the credential file to the S3 bucke
- H. Point the application to the S3 bucket.
- I. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (or each EC2 instanc
- J. Attach the new EBS volume to each EC2 instanc
- K. Migrate the credential file to the new EBS volum
- L. Point the application to the new EBS volume.

Answer: A

Explanation:

<https://aws.amazon.com/cn/blogs/security/how-to-connect-to-aws-secrets-manager-service-within-a-virtual-private-cloud/>

<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>

NEW QUESTION 47

- (Topic 1)

A company is designing an application. The application uses an AWS Lambda function to receive information through Amazon API Gateway and to store the information in an Amazon Aurora PostgreSQL database.

During the proof-of-concept stage, the company has to increase the Lambda quotas significantly to handle the high volumes of data that the company needs to load into the database. A solutions architect must recommend a new design to improve scalability and minimize the configuration effort.

Which solution will meet these requirements?

- A. Refactor the Lambda function code to Apache Tomcat code that runs on Amazon EC2 instance
- B. Connect the database by using native Java Database Connectivity (JDBC) drivers.
- C. Change the platform from Aurora to Amazon DynamoD
- D. Provision a DynamoDB Accelerator (DAX) cluste
- E. Use the DAX client SDK to point the existing DynamoDB API calls at the DAX cluster.
- F. Set up two Lambda function
- G. Configure one function to receive the informatio

- H. Configure the other function to load the information into the databases
- I. Integrate the Lambda functions by using Amazon Simple Notification Service (Amazon SNS).
- J. Set up two Lambda function
- K. Configure one function to receive the information
- L. Configure the other function to load the information into the databases
- M. Integrate the Lambda functions by using an Amazon Simple Queue Service (Amazon SQS) queue.

Answer: B

Explanation:

bottlenecks can be avoided with queues (SQS).

NEW QUESTION 51

- (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 layer
- 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 group
- 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 group
- 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-cognito/module-4/>

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 52

- (Topic 1)

A company recently launched a variety of new workloads on Amazon EC2 instances in its AWS account. The company needs to create a strategy to access and administer the instances remotely and securely. The company needs to implement a repeatable process that works with native AWS services and follows the AWS Well-Architected Framework.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use the EC2 serial console to directly access the terminal interface of each instance for administration.
- B. Attach the appropriate IAM role to each existing instance and new instance
- C. Use AWS Systems Manager Session Manager to establish a remote SSH session.
- D. Create an administrative SSH key pair
- E. Load the public key into each EC2 instance
- F. Deploy a bastion host in a public subnet to provide a tunnel for administration of each instance.
- G. Establish an AWS Site-to-Site VPN connection
- H. Instruct administrators to use their local on-premises machines to connect directly to the instances by using SSH keys across the VPN tunnel.

Answer: B

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/setup-launch-managed-instance.html>

NEW QUESTION 54

- (Topic 2)

A medical records company is hosting an application on Amazon EC2 instances. The application processes customer data files that are stored on Amazon S3. The EC2 instances are hosted in public subnets. The EC2 instances access Amazon S3 over the internet, but they do not require any other network access.

A new requirement mandates that the network traffic for file transfers take a private route and not be sent over the internet.

Which change to the network architecture should a solutions architect recommend to meet this requirement?

- A. Create a NAT gateway
- B. Configure the route table for the public subnets to send traffic to Amazon S3 through the NAT gateway.
- C. Configure the security group for the EC2 instances to restrict outbound traffic so that only traffic to the S3 prefix list is permitted.
- D. Move the EC2 instances to private subnet
- E. Create a VPC endpoint for Amazon S3, and link the endpoint to the route table for the private subnets
- F. Remove the internet gateway from the VPC
- G. Set up an AWS Direct Connect connection, and route traffic to Amazon S3 over the Direct Connect connection.

Answer: C

Explanation:

To meet the new requirement of transferring files over a private route, the EC2 instances should be moved to private subnets, which do not have direct access to the internet. This ensures that the traffic for file transfers does not go over the internet. To enable the EC2 instances to access Amazon S3, a VPC endpoint for Amazon S3 can be created. VPC endpoints allow resources within a VPC to communicate with resources in other services without the traffic being sent over the internet. By linking the VPC endpoint to the route table for the private subnets, the EC2 instances can access Amazon S3 over a private connection within the VPC.

NEW QUESTION 58

- (Topic 2)

A company needs to save the results from a medical trial to an Amazon S3 repository. The repository must allow a few scientists to add new files and must restrict all other users to read-only access. No users can have the ability to modify or delete any files in the repository. The company must keep every file in the repository for a minimum of 1 year after its creation date.

Which solution will meet these requirements?

- A. Use S3 Object Lock In governance mode with a legal hold of 1 year
- B. Use S3 Object Lock in compliance mode with a retention period of 365 days.
- C. Use an IAM role to restrict all users from deleting or changing objects in the S3 bucket Use an S3 bucket policy to only allow the IAM role
- D. Configure the S3 bucket to invoke an AWS Lambda function every time an object is added Configure the function to track the hash of the saved object to that modified objects can be marked accordingly

Answer: B

Explanation:

In compliance mode, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode helps ensure that an object version can't be overwritten or deleted for the duration of the retention period. In governance mode, users can't overwrite or delete an object version or alter its lock settings unless they have special permissions. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. In Governance mode, Objects can be deleted by some users with special permissions, this is against the requirement.

Compliance:

- Object versions can't be overwritten or deleted by any user, including the root user
- Objects retention modes can't be changed, and retention periods can't be shortened

Governance:

- Most users can't overwrite or delete an object version or alter its lock settings
- Some users have special permissions to change the retention or delete the object

NEW QUESTION 60

- (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 access
- 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 access
- 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

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_vpc.html#example_vpc_2

NEW QUESTION 63

- (Topic 2)

A company uses a popular content management system (CMS) for its corporate website. However, the required patching and maintenance are burdensome. The company is redesigning its website and wants a new solution. The website will be updated four times a year and does not need to have any dynamic content available. The solution must provide high scalability and enhanced security.

Which combination of changes will meet these requirements with the LEAST operational overhead? (Choose two.)

- A. Deploy an AWS WAF web ACL in front of the website to provide HTTPS functionality
- B. Create and deploy an AWS Lambda function to manage and serve the website content
- C. Create the new website and an Amazon S3 bucket Deploy the website on the S3 bucket with static website hosting enabled
- D. Create the new website
- E. Deploy the website by using an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer.

Answer: AD

Explanation:

A -> We can configure CloudFront to require HTTPS from clients (enhanced security)

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-https-viewers-to-cloudfront.html> D -> storing static website on S3 provides scalability and less operational overhead, then configuration of Application LB and EC2 instances (hence E is out)

NEW QUESTION 64

- (Topic 2)

A company runs an application using Amazon ECS. The application creates resized versions of an original image and then makes Amazon S3 API calls to store the resized images in Amazon S3.

How can a solutions architect ensure that the application has permission to access Amazon S3?

- A. Update the S3 role in AWS IAM to allow read/write access from Amazon ECS, and then relaunch the container.
- B. Create an IAM role with S3 permissions, and then specify that role as the taskRoleArn in the task definition.
- C. Create a security group that allows access from Amazon ECS to Amazon S3, and update the launch configuration used by the ECS cluster.
- D. Create an IAM user with S3 permissions, and then relaunch the Amazon EC2 instances for the ECS cluster while logged in as this account.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ecs-taskdefinition.html>

NEW QUESTION 67

- (Topic 2)

An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS.

The company wants to make all the data available to various teams so that the teams can perform analytics. The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead.

Which solution will meet these requirements?

- A. Migrate the purchase data to write directly to Amazon RD
- B. Use RDS access controls to limit access.
- C. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3. Create an AWS Glue crawle
- D. Use Amazon Athena to query the dat
- E. Use S3 policies to limit access.
- F. Create a data lake by using AWS Lake Formatio
- G. Create an AWS Glue JDBC connection to Amazon RD
- H. Register (he S3 bucket in Lake Formatio
- I. Use Lake Formation access controls to limit access.
- J. Create an Amazon Redshift cluste
- K. Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshif
- L. Use Amazon Redshift access controls to limit access.

Answer: C

Explanation:

To make all the data available to various teams and minimize operational overhead, the company can create a data lake by using AWS Lake Formation. This will allow the company to centralize all the data in one place and use fine-grained access controls to manage access to the data. To meet the requirements of the company, the solutions architect can create a data lake by using AWS Lake Formation, create an AWS Glue JDBC connection to Amazon RDS, and register the S3 bucket in Lake Formation. The solutions architect can then use Lake Formation access controls to limit access to the data. This solution will provide the ability to manage fine-grained permissions for the data and minimize operational overhead.

NEW QUESTION 69

- (Topic 2)

A reporting team receives files each day in an Amazon S3 bucket. The reporting team manually reviews and copies the files from this initial S3 bucket to an analysis S3 bucket each day at the same time to use with Amazon QuickSight. Additional teams are starting to send more files in larger sizes to the initial S3 bucket.

The reporting team wants to move the files automatically analysis S3 bucket as the files enter the initial S3 bucket. The reporting team also wants to use AWS Lambda functions to run pattern-matching code on the copied data. In addition, the reporting team wants to send the data files to a pipeline in Amazon SageMaker Pipelines.

What should a solutions architect do to meet these requirements with the LEAST operational overhead?

- A. Create a Lambda function to copy the files to the analysis S3 bucke
- B. Create an S3 event notification for the analysis S3 bucke
- C. Configure Lambda and SageMaker Pipelines as destinations of the event notificatio
- D. Configure s3objectCreated:Put as the event type.
- E. Create a Lambda function to copy the files to the analysis S3 bucke
- F. Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon CloudWatch Events). Configure an ObjectCreated rule in EventBridge (CloudWatch Events). Configure Lambda and SageMaker Pipelines as targets for the rule.
- G. Configure S3 replication between the S3 bucket
- H. Create an S3 event notification for the analysis S3 bucke
- I. Configure Lambda and SageMaker Pipelines as destinations of the event notificatio
- J. Configure s3objectCreated:Put as the event type.
- K. Configure S3 replication between the S3 bucket
- L. Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon CloudWatch Events). Configure an ObjectCreated rule in EventBridge (CloudWatch Events). Configure Lambda and SageMaker Pipelines as targets for the rule.

Answer: D

Explanation:

This solution meets the requirements of moving the files automatically, running Lambda functions on the copied data, and sending the data files to SageMaker Pipelines with the least operational overhead. S3 replication can copy the files from the initial S3 bucket to the analysis S3 bucket as they arrive. The analysis S3 bucket can send event notifications to Amazon EventBridge (Amazon CloudWatch Events) when an object is created. EventBridge can trigger Lambda and SageMaker Pipelines as targets for the ObjectCreated rule. Lambda can run pattern-matching code on the copied data, and SageMaker Pipelines can execute a pipeline with the data files.

Option A is incorrect because creating a Lambda function to copy the files to the analysis S3 bucket is not necessary when S3 replication can do that automatically. It also adds operational overhead to manage the Lambda function. Option B is incorrect because creating a Lambda function to copy the files to the analysis S3 bucket is not necessary when S3 replication can do that automatically. It also adds operational overhead to manage the Lambda function. Option C is incorrect because using S3 event notification with multiple destinations can result in throttling or delivery failures if there are too many events. References:

? <https://aws.amazon.com/blogs/machine-learning/automate-feature-engineering-pipelines-with-amazon-sagemaker/>

? <https://docs.aws.amazon.com/sagemaker/latest/dg/automating-sagemaker-with-eventbridge.html>

? <https://aws.amazon.com/about-aws/whats-new/2021/04/new-options-trigger-amazon-sagemaker-pipeline-executions/>

NEW QUESTION 72

- (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 in 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 Kinesis 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: BE

Explanation:

<https://aws.amazon.com/blogs/big-data/top-10-performance-tuning-tips-for-amazon-athena/>
This solution meets the requirements of measuring the effectiveness of marketing campaigns by performing batch processing on csv files of sales data and storing the results in an Amazon S3 bucket once every hour. An AWS duo ETL process can use services such as AWS Glue or AWS Data Pipeline to extract data from S3, transform it into a more efficient format such as Apache Parquet, and load it back into S3. Apache Parquet is a columnar storage format that can improve the query performance and reliability of Athena by reducing the amount of data scanned, improving compression ratio, and enabling predicate pushdown.

NEW QUESTION 77

- (Topic 2)

A company is running an online transaction processing (OLTP) workload on AWS. This workload uses an unencrypted Amazon RDS DB instance in a Multi-AZ deployment. Daily database snapshots are taken from this instance.

What should a solutions architect do to ensure the database and snapshots are always encrypted moving forward?

- A. Encrypt a copy of the latest DB snapsho
- B. Replace existing DB instance by restoring the encrypted snapshot
- C. Create a new encrypted Amazon Elastic Block Store (Amazon EBS) volume and copy the snapshots to it Enable encryption on the DB instance
- D. Copy the snapshots and enable encryption using AWS Key Management Service (AWS KMS) Restore encrypted snapshot to an existing DB instance
- E. Copy the snapshots to an Amazon S3 bucket that is encrypted using server-side encryption with AWS Key Management Service (AWS KMS) managed keys (SSE-KMS)

Answer: A

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html#USER_RestoreFromSnapshot.CON
Under "Encrypt unencrypted resources" - <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

NEW QUESTION 80

- (Topic 2)

A media company is evaluating the possibility of moving rts systems to the AWS Cloud The company needs at least 10 TB of storage with the maximum possible I/O performance for video processing. 300 TB of very durable storage for storing media content, and 900 TB of storage to meet requirements for archival media that is not in use anymore

Which set of services should a solutions architect recommend to meet these requirements?

- A. Amazon EBS for maximum performance, Amazon S3 for durable data storage, and Amazon S3 Glacier for archival storage
- B. Amazon EBS for maximum performance, Amazon EFS for durable data storage and Amazon S3 Glacier for archival storage
- C. Amazon EC2 instance store for maximum performanc
- D. Amazon EFS for durable data storage and Amazon S3 for archival storage
- E. Amazon EC2 Instance store for maximum performanc
- F. Amazon S3 for durable data storage, and Amazon S3 Glacier for archival storage

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html>

NEW QUESTION 82

- (Topic 2)

A company runs an Oracle database on premises. As part of the company's migration to AWS, the company wants to upgrade the database to the most recent available version. The company also wants to set up disaster recovery (DR) for the database. The company needs to minimize the operational overhead for normal operations and DR setup. The company also needs to maintain access to the database's underlying operating system.

Which solution will meet these requirements?

- A. Migrate the Oracle database to an Amazon EC2 instanc
- B. Set up database replication to a different AWS Region.
- C. Migrate the Oracle database to Amazon RDS for Oracl
- D. Activate Cross-Region automated backups to replicate the snapshots to another AWS Region.
- E. Migrate the Oracle database to Amazon RDS Custom for Oracl
- F. Create a read replica for the database in another AWS Region.
- G. Migrate the Oracle database to Amazon RDS for Oracl
- H. Create a standby database in another Availability Zone.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/rds-custom.html> and <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/working-with-custom-oracle.html>

NEW QUESTION 83

- (Topic 2)

A solutions architect is optimizing a website for an upcoming musical event. Videos of the performances will be streamed in real time and then will be available on demand. The event is expected to attract a global online audience.

Which service will improve the performance of both the real-time and on-demand streaming?

- A. Amazon CloudFront
- B. AWS Global Accelerator
- C. Amazon Route 53
- D. Amazon S3 Transfer Acceleration

Answer: A

Explanation:

You can use CloudFront to deliver video on demand (VOD) or live streaming video using any HTTP origin. One way you can set up video workflows in the cloud is by using CloudFront together with AWS Media Services. <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/on-demand-streaming-video.html>

NEW QUESTION 84

- (Topic 2)

A company wants to run a gaming application on Amazon EC2 instances that are part of an Auto Scaling group in the AWS Cloud. The application will transmit data by using UDP packets. The company wants to ensure that the application can scale out and in as traffic increases and decreases.

What should a solutions architect do to meet these requirements?

- A. Attach a Network Load Balancer to the Auto Scaling group
- B. Attach an Application Load Balancer to the Auto Scaling group.
- C. Deploy an Amazon Route 53 record set with a weighted policy to route traffic appropriately
- D. Deploy a NAT instance that is configured with port forwarding to the EC2 instances in the Auto Scaling group.

Answer: A

Explanation:

This solution meets the requirements of running a gaming application that transmits data by using UDP packets and scaling out and in as traffic increases and decreases. A Network Load Balancer can handle millions of requests per second while maintaining high throughput at ultra low latency, and it supports both TCP and UDP protocols. An Auto Scaling group can automatically adjust the number of EC2 instances based on the demand and the scaling policies.

Option B is incorrect because an Application Load Balancer does not support UDP protocol, only HTTP and HTTPS. Option C is incorrect because Amazon Route 53 is a DNS service that can route traffic based on different policies, but it does not provide load balancing or scaling capabilities. Option D is incorrect because a NAT instance is used to enable instances in a private subnet to connect to the internet or other AWS services, but it does not provide load balancing or scaling capabilities.

References:

? <https://aws.amazon.com/blogs/aws/new-udp-load-balancing-for-network-load-balancer/>

? <https://docs.aws.amazon.com/autoscaling/ec2/userguide/AutoScalingGroup.html>

NEW QUESTION 85

- (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 user
- C. Grant the user read permission to objects in the S3 bucket
- 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 distribution
- 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.html#private-content-restricting-access-to-s3-overview>

NEW QUESTION 88

- (Topic 2)

A business's backup data totals 700 terabytes (TB) and is kept in network attached storage (NAS) at its data center. This backup data must be available in the event of occasional regulatory inquiries and preserved for a period of seven years. The organization has chosen to relocate its backup data from its on-premises data center to Amazon Web Services (AWS). Within one month, the migration must be completed. The company's public internet connection provides 500 Mbps of dedicated capacity for data transport.

What should a solutions architect do to ensure that data is migrated and stored at the LOWEST possible cost?

- A. Order AWS Snowball devices to transfer the data
- B. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
- C. Deploy a VPN connection between the data center and Amazon VPC
- D. Use the AWS CLI to copy the data from on premises to Amazon S3 Glacier.
- E. Provision a 500 Mbps AWS Direct Connect connection and transfer the data to Amazon S3. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
- F. Use AWS DataSync to transfer the data and deploy a DataSync agent on premise
- G. Use the DataSync task to copy files from the on-premises NAS storage to Amazon S3 Glacier.

Answer: A

Explanation:

<https://www.omnicalculator.com/other/data-transfer>

NEW QUESTION 92

- (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 97

- (Topic 2)

A company's application is having performance issues. The application is stateful and needs to complete m-memory tasks on Amazon EC2 instances. The company used AWS CloudFormation to deploy infrastructure and used the M5 EC2 Instance family. As traffic increased, the application performance degraded. Users are reporting delays when they attempt to access the application.

Which solution will resolve these issues in the MOST operationally efficient way?

- A. Replace the EC2 instances with T3 EC2 instances that run in an Auto Scaling group
- B. Make the changes by using the AWS Management Console.
- C. Modify the CloudFormation templates to run the EC2 instances in an Auto Scaling group
- D. Increase the desired capacity and the maximum capacity of the Auto Scaling group manually when an increase is necessary
- E. Modify the CloudFormation template
- F. Replace the EC2 instances with R5 EC2 instances
- G. Use Amazon CloudWatch built-in EC2 memory metrics to track the application performance for future capacity planning.
- H. Modify the CloudFormation template
- I. Replace the EC2 instances with R5 EC2 instances
- J. Deploy the Amazon CloudWatch agent on the EC2 instances to generate custom application latency metrics for future capacity planning.

Answer: D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudwatch-memory-metrics-ec2/>

NEW QUESTION 102

- (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 Amplify
- C. Connect the application to an Amazon API Gateway API that is integrated with AWS Lambda.
- D. Host the application on Amazon EC2 instances
- 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 106

- (Topic 2)

A solutions architect is designing a customer-facing application for a company. The application's database will have a clearly defined access pattern throughout the year and will have a variable number of reads and writes that depend on the time of year. The company must retain audit records for the database for 7 days. The recovery point objective (RPO) must be less than 5 hours.

Which solution meets these requirements?

- A. Use Amazon DynamoDB with auto scaling. Use on-demand backups and Amazon DynamoDB Streams
- B. Use Amazon Redshift

- C. Configure concurrency scalin
- D. Activate audit loggin
- E. Perform database snapshots every 4 hours.
- F. Use Amazon RDS with Provisioned IOPS Activate the database auditing parameter Perform database snapshots every 5 hours
- G. Use Amazon Aurora MySQL with auto scalin
- H. Activate the database auditing parameter

Answer: A

Explanation:

This solution meets the requirements of a customer-facing application that has a clearly defined access pattern throughout the year and a variable number of reads and writes that depend on the time of year. Amazon DynamoDB is a fully managed NoSQL database service that can handle any level of request traffic and data size. DynamoDB auto scaling can automatically adjust the provisioned read and write capacity based on the actual workload. DynamoDB on-demand backups can create full backups of the tables for data protection and archival purposes. DynamoDB Streams can capture a time-ordered sequence of item-level modifications in the tables for audit purposes.

Option B is incorrect because Amazon Redshift is a data warehouse service that is designed for analytical workloads, not for customer-facing applications. Option C is incorrect because Amazon RDS with Provisioned IOPS can provide consistent performance for relational databases, but it may not be able to handle unpredictable spikes in traffic and data size. Option D is incorrect because Amazon Aurora MySQL with auto scaling can provide high performance and availability for relational databases, but it does not support audit logging as a parameter.

References:

- ? <https://aws.amazon.com/dynamodb/>
- ? <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScalin g.html>
- ? <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/BackupRe store.html>
- ? <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.ht ml>

NEW QUESTION 107

- (Topic 2)

A company is running a multi-tier web application on premises. The web application is containerized and runs on a number of Linux hosts connected to a PostgreSQL database that contains user records. The operational overhead of maintaining the infrastructure and capacity planning is limiting the company's growth. A solutions architect must improve the application's infrastructure.

Which combination of actions should the solutions architect take to accomplish this? (Choose two.)

- A. Migrate the PostgreSQL database to Amazon Aurora
- B. Migrate the web application to be hosted on Amazon EC2 instances.
- C. Set up an Amazon CloudFront distribution for the web application content.
- D. Set up Amazon ElastiCache between the web application and the PostgreSQL database.
- E. Migrate the web application to be hosted on AWS Fargate with Amazon Elastic Container Service (Amazon ECS).

Answer: AE

Explanation:

Amazon Aurora is a fully managed, scalable, and highly available relational database service that is compatible with PostgreSQL. Migrating the database to Amazon Aurora would reduce the operational overhead of maintaining the database infrastructure and allow the company to focus on building and scaling the application. AWS Fargate is a fully managed container orchestration service that enables users to run containers without the need to manage the underlying EC2 instances. By using AWS Fargate with Amazon Elastic Container Service (Amazon ECS), the solutions architect can improve the scalability and efficiency of the web application and reduce the operational overhead of maintaining the underlying infrastructure.

NEW QUESTION 109

- (Topic 3)

A company has migrated an application to Amazon EC2 Linux instances. One of these EC2 instances runs several 1-hour tasks on a schedule. These tasks were written by different teams and have no common programming language. The company is concerned about performance and scalability while these tasks run on a single instance. A solutions architect needs to implement a solution to resolve these concerns.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Batch to run the tasks as job
- B. Schedule the jobs by using Amazon EventBridge (Amazon CloudWatch Events).
- C. Convert the EC2 instance to a containe
- D. Use AWS App Runner to create the container on demand to run the tasks as jobs.
- E. Copy the tasks into AWS Lambda function
- F. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events).
- G. Create an Amazon Machine Image (AMI) of the EC2 instance that runs the task
- H. Create an Auto Scaling group with the AMI to run multiple copies of the instance.

Answer: A

Explanation:

AWS Batch is a fully managed service that enables users to run batch jobs on AWS. It can handle different types of tasks written in different languages and run them on EC2 instances. It also integrates with Amazon EventBridge (Amazon CloudWatch Events) to schedule jobs based on time or event triggers. This solution will meet the requirements of performance, scalability and low operational overhead¹².

* B. Convert the EC2 instance to a container. Use AWS App Runner to create the container on demand to run the tasks as jobs. This solution will not meet the requirement of low operational overhead, as it involves converting the EC2 instance to a container and using AWS App Runner, which is a service that automatically builds and deploys web applications and load balances traffic². This is not necessary for running batch jobs.

* C. Copy the tasks into AWS Lambda functions. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events). This solution will not meet the requirement of performance, as AWS Lambda has a limit of 15 minutes for execution time and 10 GB for memory allocation³. These limits may not be sufficient for running 1-hour tasks.

* D. Create an Amazon Machine Image (AMI) of the EC2 instance that runs the tasks. Create an Auto Scaling group with the AMI to run multiple copies of the instance. This solution will not meet the requirement of low operational overhead, as it involves creating and maintaining AMIs and Auto Scaling groups, which are additional resources that need to be configured and managed².

Reference URL: <https://docs.aws.amazon.com/whitepapers/latest/aws-overview/compute- services.html>

NEW QUESTION 112

- (Topic 3)

A gaming company is moving its public scoreboard from a data center to the AWS Cloud. The company uses Amazon EC2 Windows Server instances behind an Application Load Balancer to host its dynamic application. The company needs a highly available storage solution for the application. The application consists of static files and dynamic server-side code.

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

- A. Store the static files on Amazon S3. Use Amazon CloudFront to cache objects at the edge.
- B. Store the static files on Amazon S3. Use Amazon ElastiCache to cache objects at the edge.
- C. Store the server-side code on Amazon Elastic File System (Amazon EFS). Mount the EFS volume on each EC2 instance to share the files.
- D. Store the server-side code on Amazon FSx for Windows File Server
- E. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.
- F. Store the server-side code on a General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volume
- G. Mount the EBS volume on each EC2 instance to share the files.

Answer: AD

Explanation:

A because ElastiCache, despite being ideal for leaderboards per Amazon, doesn't cache at edge locations. D because FSx has higher performance for low latency needs. <https://www.techtarget.com/searchaws/tip/Amazon-FSx-vs-EFS-Compare-the-AWS-file-services> "FSx is built for high performance and submillisecond latency using solid-state drive storage volumes. This design enables users to select storage capacity and latency independently. Thus, even a subterabyte file system can have 256 Mbps or higher throughput and support volumes up to 64 TB."

Amazon S3 is an object storage service that can store static files such as images, videos, documents, etc. Amazon EFS is a file storage service that can store files in a hierarchical structure and supports NFS protocol. Amazon FSx for Windows File Server is a file storage service that can store files in a hierarchical structure and supports SMB protocol. Amazon EBS is a block storage service that can store data in fixed-size blocks and attach to EC2 instances.

Based on these definitions, the combination of steps that should be taken to meet the requirements are:

* A. Store the static files on Amazon S3. Use Amazon CloudFront to cache objects at the edge. D. Store the server-side code on Amazon FSx for Windows File Server. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.

NEW QUESTION 113

- (Topic 3)

A company wants to migrate a Windows-based application from on premises to the AWS Cloud. The application has three tiers, a business tier, and a database tier with Microsoft SQL Server. The company wants to use specific features of SQL Server such as native backups and Data Quality Services. The company also needs to share files for process between the tiers.

How should a solution architect design the architecture to meet these requirements?

- A. Host all three on Amazon instance
- B. Use Amazon FSx File Gateway for file sharing between tiers.
- C. Host all three on Amazon EC2 instance
- D. Use Amazon FSx for Windows file sharing between the tiers.
- E. Host the application tier and the business tier on Amazon EC2 instance
- F. Host the database tier on Amazon RD
- G. Use Amazon Elastic File system (Amazon EFS) for file sharing between the tiers.
- H. Host the application tier and the business tier on Amazon EC2 instance
- I. Host the database tier on Amazon RD
- J. Use a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume for file sharing between the tiers.

Answer: B

Explanation:

This solution will allow the company to host all three tiers on Amazon EC2 instances while using Amazon FSx for Windows File Server to provide Windows-based file sharing between the tiers. This will allow the company to use specific features of SQL Server, such as native backups and Data Quality Services, while sharing files for processing between the tiers.

NEW QUESTION 116

- (Topic 3)

A company is building a new web-based customer relationship management application. The application will use several Amazon EC2 instances that are backed by Amazon Elastic Block Store (Amazon EBS) volumes behind an Application Load Balancer (ALB). The application will also use an Amazon Aurora database. All data for the application must be encrypted at rest and in transit.

Which solution will meet these requirements?

- A. Use AWS Key Management Service (AWS KMS) certificates on the ALB to encrypt data in transit
- B. Use AWS Certificate Manager (ACM) to encrypt the EBS volumes and Aurora database storage at rest.
- C. Use the AWS root account to log in to the AWS Management Console
- D. Upload the company's encryption certificate
- E. While in the root account, select the option to turn on encryption for all data at rest and in transit for the account.
- F. Use a AWS Key Management Service (AWS KMS) to encrypt the EBS volumes and Aurora database storage at rest
- G. Attach an AWS Certificate Manager (ACM) certificate to the ALB to encrypt data in transit.
- H. Use BitLocker to encrypt all data at rest
- I. Import the company's TLS certificate keys to AWS Key Management Service (AWS KMS). Attach the KMS keys to the ALB to encrypt data in transit.

Answer: C

Explanation:

This option is the most efficient because it uses AWS Key Management Service (AWS KMS), which is a service that makes it easy for you to create and manage cryptographic keys and control their use across a wide range of AWS services and with your applications running on AWS. It also uses AWS KMS to encrypt the EBS volumes and Aurora database storage at rest, which provides data protection by encrypting your data with encryption keys that you manage. It also uses AWS Certificate Manager (ACM), which is a service that lets you easily provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and your internal connected resources. It also attaches an ACM certificate to the ALB to encrypt data in transit, which provides data protection by enabling SSL/TLS encryption for connections between clients and the load balancer. This solution meets the requirement of encrypting all data for the application at rest and in transit. Option A is less efficient because it uses AWS KMS certificates on the ALB to encrypt data in transit,

which is not possible as AWS KMS does not provide certificates but only keys. It also uses AWS Certificate Manager (ACM) to encrypt the EBS volumes and Aurora database storage at rest, which is not possible as ACM does not provide encryption but only certificates. Option B is less efficient because it uses the AWS root account to log in to the AWS Management Console, which is not recommended as it has unrestricted access to all resources in your account. It also uploads the company's encryption certificates, which is not necessary as ACM can provide certificates for free. It also selects the option to turn on encryption for all data at rest and in transit for the account, which is not possible as encryption settings are specific to each service and resource. Option D is less efficient because it uses BitLocker to encrypt all data at rest, which is a Windows feature that provides encryption for volumes on Windows servers. However, this does not provide encryption for Aurora database storage at rest, as Aurora runs on Linux servers. It also imports the company's TLS certificate keys to AWS KMS, which is not necessary as ACM can provide certificates for free. It also attaches the KMS keys to the ALB to encrypt data in transit, which is not possible as ALB requires certificates and not keys.

NEW QUESTION 119

- (Topic 3)

A company's facility has badge readers at every entrance throughout the building. When badges are scanned, the readers send a message over HTTPS to indicate who attempted to access that particular entrance.

A solutions architect must design a system to process these messages from the sensors. The solution must be highly available, and the results must be made available for the company's security team to analyze.

Which system architecture should the solutions architect recommend?

- A. Launch an Amazon EC2 instance to serve as the HTTPS endpoint and to process the messages Configure the EC2 instance to save the results to an Amazon S3 bucket.
- B. Create an HTTPS endpoint in Amazon API Gateway
- C. Configure the API Gateway endpoint to invoke an AWS Lambda function to process the messages and save the results to an Amazon DynamoDB table.
- D. Use Amazon Route 53 to direct incoming sensor messages to an AWS Lambda function
- E. Configure the Lambda function to process the messages and save the results to an Amazon DynamoDB table.
- F. Create a gateway VPC endpoint for Amazon S3. Configure a Site-to-Site VPN connection from the facility network to the VPC so that sensor data can be written directly to an S3 bucket by way of the VPC endpoint.

Answer: B

Explanation:

Deploy Amazon API Gateway as an HTTPS endpoint and AWS Lambda to process and save the messages to an Amazon DynamoDB table. This option provides a highly available and scalable solution that can easily handle large amounts of data. It also integrates with other AWS services, making it easier to analyze and visualize the data for the security team.

NEW QUESTION 124

- (Topic 3)

A company needs to migrate a legacy application from an on-premises data center to the AWS Cloud because of hardware capacity constraints. The application runs 24 hours a day, 7 days a week. The application database storage continues to grow over time.

What should a solution architect do to meet these requirements MOST cost-effectively?

- A. Migrate the application layer to Amazon EC2 Spot Instances Migrate the data storage layer to Amazon S3.
- B. Migrate the application layer to Amazon EC2 Reserved Instances Migrate the data storage layer to Amazon RDS On-Demand Instances.
- C. Migrate the application layer to Amazon EC2 Reserved instances Migrate the data storage layer to Amazon Aurora Reserved Instances.
- D. Migrate the application layer to Amazon EC2 On Demand Amazon Migrate the data storage layer to Amazon RDS Reserved instances.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.AuroraMySQL.html>

NEW QUESTION 129

- (Topic 3)

A rapidly growing global ecommerce company is hosting its web application on AWS. The web application includes static content and dynamic content. The website stores online transaction processing (OLTP) data in an Amazon RDS database. The website's users are experiencing slow page loads.

Which combination of actions should a solutions architect take to resolve this issue? (Select TWO.)

- A. Configure an Amazon Redshift cluster.
- B. Set up an Amazon CloudFront distribution
- C. Host the dynamic web content in Amazon S3
- D. Create a read replica for the RDS DB instance.
- E. Configure a Multi-AZ deployment for the RDS DB instance

Answer: BD

Explanation:

To resolve the issue of slow page loads for a rapidly growing e-commerce website hosted on AWS, a solutions architect can take the following two actions:

- * 1. Set up an Amazon CloudFront distribution
- * 2. Create a read replica for the RDS DB instance

Configuring an Amazon Redshift cluster is not relevant to this issue since Redshift is a data warehousing service and is typically used for the analytical processing of large amounts of data.

Hosting the dynamic web content in Amazon S3 may not necessarily improve performance since S3 is an object storage service, not a web application server. While S3 can be used to host static web content, it may not be suitable for hosting dynamic web content since S3 doesn't support server-side scripting or processing.

Configuring a Multi-AZ deployment for the RDS DB instance will improve high availability but may not necessarily improve performance.

NEW QUESTION 134

- (Topic 3)

A company wants to deploy a new public web application on AWS The application includes a web server tier that uses Amazon EC2 instances The application also includes a database tier that uses an Amazon RDS for MySQL DB instance

The application must be secure and accessible for global customers that have dynamic IP addresses
How should a solutions architect configure the security groups to meet these requirements'?

- A. Configure the security group for the web servers to allow inbound traffic on port 443 from 0.0.0.0/0. Configure the security group for the DB instance to allow inbound traffic on port 3306 from the security group of the web servers
- B. Configure the security group for the web servers to allow inbound traffic on port 443 from the IP addresses of the customers. Configure the security group for the DB instance to allow inbound traffic on port 3306 from the security group of the web servers
- C. Configure the security group for the web servers to allow inbound traffic on port 443 from the IP addresses of the customers. Configure the security group for the DB instance to allow inbound traffic on port 3306 from the IP addresses of the customers
- D. Configure the security group for the web servers to allow inbound traffic on port 443 from 0.0.0.0/0. Configure the security group for the DB instance to allow inbound traffic on port 3306 from 0.0.0.0/0

Answer: A

Explanation:

? Restricting inbound access to the web servers to only port 443, which is used for HTTPS traffic, and allowing access from any IP address (0.0.0.0/0), since the application is public and accessible for global customers.

? Restricting inbound access to the DB instance to only port 3306, which is used for MySQL traffic, and allowing access only from the security group of the web servers, which creates a secure connection between the two tiers and prevents unauthorized access to the database.

? Restricting outbound access to the minimum required for both tiers, which is not specified in the question but can be assumed to be similar to the inbound rules.

References:

? Security groups - Amazon Virtual Private Cloud

? 5 Best Practices for AWS Security Groups - DZone

NEW QUESTION 135

- (Topic 3)

A company is using AWS to design a web application that will process insurance quotes. Users will request quotes from the application. Quotes must be separated by quote type, must be responded to within 24 hours, and must not get lost. The solution must maximize operational efficiency and must minimize maintenance. Which solution meets these requirements?

- A. Create multiple Amazon Kinesis data streams based on the quote type. Configure the web application to send messages to the proper data stream. Configure each backend group of application servers to use the Kinesis Client Library (KCL) to pool messages from its own data stream.
- B. Create an AWS Lambda function and an Amazon Simple Notification Service (Amazon SNS) topic for each quote type. Subscribe the Lambda function to its associated SNS topic. Configure the application to publish requests for quotes to the appropriate SNS topic.
- C. Create a single Amazon Simple Notification Service (Amazon SNS) topic. Subscribe Amazon Simple Queue Service (Amazon SQS) queues to the SNS topic. Configure SNS message filtering to publish messages to the proper SQS queue based on the quote type. Configure each backend application server to use its own SQS queue.
- D. Create multiple Amazon Kinesis Data Firehose delivery streams based on the quote type to deliver data streams to an Amazon Elasticsearch Service (Amazon ES) cluster. Configure the application to send messages to the proper delivery stream. Configure each backend group of application servers to search for the messages from Amazon ES and process them accordingly.

Answer: C

Explanation:

<https://aws.amazon.com/getting-started/hands-on/filter-messages-published-to-topics/>

NEW QUESTION 136

- (Topic 3)

A company is deploying a two-tier web application in a VPC. The web tier is using an Amazon EC2 Auto Scaling group with public subnets that span multiple Availability Zones. The database tier consists of an Amazon RDS for MySQL DB instance in separate private subnets. The web tier requires access to the database to retrieve product information.

The web application is not working as intended. The web application reports that it cannot connect to the database. The database is confirmed to be up and running. All configurations for the network ACLs, security groups, and route tables are still in their default states.

What should a solutions architect recommend to fix the application?

- A. Add an explicit rule to the private subnet's network ACL to allow traffic from the web tier's EC2 instances.
- B. Add a route in the VPC route table to allow traffic between the web tier's EC2 instances and the database tier.
- C. Deploy the web tier's EC2 instances and the database tier's RDS instance into two separate VPCs.
- D. Add and configure VPC peering.
- E. Add an inbound rule to the security group of the database tier's RDS instance to allow traffic from the web tier's security group.

Answer: D

Explanation:

This answer is correct because it allows the web tier to access the database tier by using security groups as a source, which is a recommended best practice for VPC connectivity. Security groups are stateful and can reference other security groups in the same VPC, which simplifies the configuration and maintenance of the firewall rules. By adding an inbound rule to the database tier's security group, the web tier's EC2 instances can connect to the RDS instance on port 3306, regardless of their IP addresses or subnets. References:

? Security groups - Amazon Virtual Private Cloud

? Best practices and reference architectures for VPC design

NEW QUESTION 140

- (Topic 3)

A company must migrate 20 TB of data from a data center to the AWS Cloud within 30 days. The company's network bandwidth is limited to 15 Mbps and cannot exceed 70% utilization. What should a solutions architect do to meet these requirements?

- A. Use AWS Snowball.
- B. Use AWS DataSync.
- C. Use a secure VPN connection.
- D. Use Amazon S3 Transfer Acceleration.

Answer: A

Explanation:

AWS Snowball is a secure data transport solution that accelerates moving large amounts of data into and out of the AWS cloud. It can move up to 80 TB of data at a time, and provides a network bandwidth of up to 50 Mbps, so it is well-suited for the task. Additionally, it is secure and easy to use, making it the ideal solution for this migration.

NEW QUESTION 143

- (Topic 3)

A rapidly growing ecommerce company is running its workloads in a single AWS Region. A solutions architect must create a disaster recovery (DR) strategy that includes a different AWS Region. The company wants its database to be up to date in the DR Region with the least possible latency. The remaining infrastructure in the DR Region needs to run at reduced capacity and must be able to scale up if necessary. Which solution will meet these requirements with the LOWEST recovery time objective (RTO)?

- A. Use an Amazon Aurora global database with a pilot light deployment
- B. Use an Amazon Aurora global database with a warm standby deployment
- C. Use an Amazon RDS Multi-AZ DB instance with a pilot light deployment
- D. Use an Amazon RDS Multi-AZ DB instance with a warm standby deployment

Answer: B

Explanation:

<https://docs.aws.amazon.com/whitepapers/latest/disaster-recovery-workloads-on-aws/disaster-recovery-options-in-the-cloud.html>

NEW QUESTION 146

- (Topic 3)

An application runs on Amazon EC2 instances in private subnets. The application needs to access an Amazon DynamoDB table. What is the MOST secure way to access the table while ensuring that the traffic does not leave the AWS network?

- A. Use a VPC endpoint for DynamoDB.
- B. Use a NAT gateway in a public subnet.
- C. Use a NAT instance in a private subnet.
- D. Use the internet gateway attached to the VPC.

Answer: A

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>

A VPC endpoint for DynamoDB enables Amazon EC2 instances in your VPC to use their private IP addresses to access DynamoDB with no exposure to the public internet. Your EC2 instances do not require public IP addresses, and you don't need an internet gateway, a NAT device, or a virtual private gateway in your VPC. You use endpoint policies to control access to DynamoDB. Traffic between your VPC and the AWS service does not leave the Amazon network.

NEW QUESTION 147

- (Topic 3)

A solutions architect is creating a new VPC design. There are two public subnets for the load balancer, two private subnets for web servers, and two private subnets for MySQL. The web servers use only HTTPS. The solutions architect has already created a security group for the load balancer allowing port 443 from 0.0.0.0/0. Company policy requires that each resource has the least access required to still be able to perform its tasks. Which additional configuration strategy should the solutions architect use to meet these requirements?

- A. Create a security group for the web servers and allow port 443 from 0.0.0.0/0. Create a security group for the MySQL servers and allow port 3306 from the web servers security group.
- B. Create a network ACL for the web servers and allow port 443 from 0.0.0.0/0. Create a network ACL for the MySQL servers and allow port 3306 from the web servers security group.
- C. Create a security group for the web servers and allow port 443 from the load balancer. Create a security group for the MySQL servers and allow port 3306 from the web servers security group.
- D. Create a network ACL for the web servers and allow port 443 from the load balancer. Create a network ACL for the MySQL servers and allow port 3306 from the web servers security group.

Answer: C

Explanation:

This answer is correct because it provides a resilient and durable replacement for the on-premises file share that is compatible with Windows IIS web servers. Amazon FSx for Windows File Server is a fully managed service that provides shared file storage built on Windows Server. It supports the SMB protocol and integrates with Microsoft Active Directory, which enables seamless access and authentication for Windows-based applications. Amazon FSx for Windows File Server also offers the following benefits:

? Resilience: Amazon FSx for Windows File Server can be deployed in multiple

Availability Zones, which provides high availability and failover protection. It also supports automatic backups and restores, as well as self-healing features that detect and correct issues.

? Durability: Amazon FSx for Windows File Server replicates data within and across Availability Zones, and stores data on highly durable storage devices. It also supports encryption at rest and in transit, as well as file access auditing and data deduplication.

? Performance: Amazon FSx for Windows File Server delivers consistent sub-millisecond latencies and high throughput for file operations. It also supports SSD storage, native Windows features such as Distributed File System (DFS) Namespaces and Replication, and user-driven performance scaling.

By configuring the Amazon FSx file share to use an AWS KMS CMK to encrypt the images in the file share, the company can protect the images from unauthorized access and comply with company policy. By using NTFS permission sets on the images, the company can prevent accidental deletion of the images by restricting who can modify or delete them. References:

? Amazon FSx for Windows File Server

? Using Microsoft Windows file shares

NEW QUESTION 151

- (Topic 3)

A meteorological startup company has a custom web application to sell weather data to its users online. The company uses Amazon DynamoDB to store its data and wants to build a new service that sends an alert to the managers of four internal teams every time a new weather event is recorded. The company does not want the new service to affect the performance of the current application.

What should a solutions architect do to meet these requirements with the LEAST amount of operational overhead?

- A. Use DynamoDB transactions to write new event data to the table. Configure the transactions to notify internal teams.
- B. Have the current application publish a message to four Amazon Simple Notification Service (Amazon SNS) topics.
- C. Have each team subscribe to one topic.
- D. Enable Amazon DynamoDB Streams on the table.
- E. Use triggers to write to a single Amazon Simple Notification Service (Amazon SNS) topic to which the teams can subscribe.
- F. Add a custom attribute to each record to flag new items.
- G. Write a cron job that scans the table every minute for items that are new and notifies an Amazon Simple Queue Service (Amazon SQS) queue to which the teams can subscribe.

Answer: C

Explanation:

The best solution to meet these requirements with the least amount of operational overhead is to enable Amazon DynamoDB Streams on the table and use triggers to write to a single Amazon Simple Notification Service (Amazon SNS) topic to which the teams can subscribe. This solution requires minimal configuration and infrastructure setup, and Amazon DynamoDB Streams provide a low-latency way to capture changes to the DynamoDB table. The triggers automatically capture the changes and publish them to the SNS topic, which notifies the internal teams.

NEW QUESTION 155

- (Topic 3)

A company has an on-premises MySQL database used by the global sales team with infrequent access patterns. The sales team requires the database to have minimal downtime. A database administrator wants to migrate this database to AWS without selecting a particular instance type in anticipation of more users in the future.

Which service should a solutions architect recommend?

- A. Amazon Aurora MySQL
- B. Amazon Aurora Serverless for MySQL
- C. Amazon Redshift Spectrum
- D. Amazon RDS for MySQL

Answer: B

Explanation:

Amazon Aurora Serverless for MySQL is a fully managed, auto-scaling relational database service that scales up or down automatically based on the application demand. This service provides all the capabilities of Amazon Aurora, such as high availability, durability, and security, without requiring the customer to provision any database instances. With Amazon Aurora Serverless for MySQL, the sales team can enjoy minimal downtime since the database is designed to automatically scale to accommodate the increased traffic. Additionally, the service allows the customer to pay only for the capacity used, making it cost-effective for infrequent access patterns. Amazon RDS for MySQL could also be an option, but it requires the customer to select an instance type, and the database administrator would need to monitor and adjust the instance size manually to accommodate the increasing traffic.

NEW QUESTION 160

- (Topic 3)

A company has an application that collects data from IoT sensors on automobiles. The data is streamed and stored in Amazon S3 through Amazon Kinesis Data Firehose. The data produces trillions of S3 objects each year. Each morning, the company uses the data from the previous 30 days to retrain a suite of machine learning (ML) models.

Four times each year, the company uses the data from the previous 12 months to perform analysis and train other ML models. The data must be available with minimal delay for up to 1 year. After 1 year, the data must be retained for archival purposes.

Which storage solution meets these requirements MOST cost-effectively?

- A. Use the S3 Intelligent-Tiering storage class.
- B. Create an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 1 year.
- C. Use the S3 Intelligent-Tiering storage class.
- D. Configure S3 Intelligent-Tiering to automatically move objects to S3 Glacier Deep Archive after 1 year.
- E. Use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.
- F. Create an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 1 year.
- G. Use the S3 Standard storage class.
- H. Create an S3 Lifecycle policy to transition objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days, and then to S3 Glacier Deep Archive after 1 year.

Answer: D

Explanation:

- First 30 days - data access every morning (predictable and frequently) – S3 Standard - After 30 days, accessed 4 times a year – S3 Infrequent Access - Data preserved - S3 Glacier Deep Archive

NEW QUESTION 165

- (Topic 3)

A company is launching a new application deployed on an Amazon Elastic Container Service (Amazon ECS) cluster and is using the Fargate launch type for ECS tasks. The company is monitoring CPU and memory usage because it is expecting high traffic to the application upon its launch. However, the company wants to reduce costs when utilization decreases.

What should a solutions architect recommend?

- A. Use Amazon EC2 Auto Scaling to scale at certain periods based on previous traffic patterns.
- B. Use an AWS Lambda function to scale Amazon ECS based on metric breaches that trigger an Amazon CloudWatch alarm.
- C. Use Amazon EC2 Auto Scaling with simple scaling policies to scale when ECS metric breaches trigger an Amazon CloudWatch alarm.

D. Use AWS Application Auto Scaling with target tracking policies to scale when ECS metric breaches trigger an Amazon CloudWatch alarm

Answer: D

Explanation:

<https://docs.aws.amazon.com/autoscaling/application/userguide/what-is-application-auto-scaling.html>

NEW QUESTION 170

- (Topic 3)

A company wants to migrate its 1 PB on-premises image repository to AWS. The images will be used by a serverless web application. Images stored in the repository are rarely accessed, but they must be immediately available. Additionally, the images must be encrypted at rest and protected from accidental deletion. Which solution meets these requirements?

- A. Implement client-side encryption and store the images in an Amazon S3 Glacier vault. Set a vault lock to prevent accidental deletion.
- B. Store the images in an Amazon S3 bucket in the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Enable versioning, default encryption, and MFA Delete on the S3 bucket.
- C. Store the images in an Amazon FSx for Windows File Server file share. Configure the Amazon FSx file share to use an AWS Key Management Service (AWS KMS) customer master key (CMK) to encrypt the images in the file share. Use NTFS permission sets on the images to prevent accidental deletion.
- D. Store the images in an Amazon Elastic File System (Amazon EFS) file share in the Infrequent Access storage class. Configure the EFS file share to use an AWS Key Management Service (AWS KMS) customer master key (CMK) to encrypt the images in the file share.
- E. Use NFS permission sets on the images to prevent accidental deletion.

Answer: B

Explanation:

This answer is correct because it provides a resilient and durable replacement for the on-premises file share that is compatible with a serverless web application. Amazon S3 is a fully managed object storage service that can store any amount of data and serve it over the internet. It supports the following features:

- ? Resilience: Amazon S3 stores data across multiple Availability Zones within a Region, and offers 99.999999999% (11 9's) of durability. It also supports cross-region replication, which enables automatic and asynchronous copying of objects across buckets in different AWS Regions.
- ? Durability: Amazon S3 encrypts data at rest using server-side encryption with either Amazon S3-managed keys (SSE-S3), AWS KMS keys (SSE-KMS), or customer-provided keys (SSE-C). It also supports encryption in transit using SSL/TLS. Amazon S3 also provides data protection features such as versioning, which keeps multiple versions of an object in the same bucket, and MFA Delete, which requires additional authentication for deleting an object version or changing the versioning state of a bucket.

? Performance: Amazon S3 delivers high performance and scalability for serving static and dynamic web content. It also supports features such as S3 Transfer Acceleration, which speeds up data transfers by routing requests to AWS edge locations, and S3 Select, which enables retrieving only a subset of data from an object by using simple SQL expressions.

The S3 Standard-Infrequent Access (S3 Standard-IA) storage class is suitable for storing images that are rarely accessed, but must be immediately available when needed. It offers the same high durability, throughput, and low latency as S3 Standard, but with a lower storage cost per GB and a higher per-request cost.

References:

? Amazon Simple Storage Service

? Storage classes - Amazon Simple Storage Service

NEW QUESTION 175

- (Topic 3)

A development team has launched a new application that is hosted on Amazon EC2 instances inside a development VPC. A solution architect needs to create a new VPC in the same account. The new VPC will be peered with the development VPC. The VPC CIDR block for the development VPC is 192.168.0.0/24. The solution architect needs to create a CIDR block for the new VPC. The CIDR block must be valid for a VPC peering connection to the development VPC.

What is the SMALLEST CIDR block that meets these requirements?

- A. 10.0.1.0/32
- B. 192.168.0.0/24
- C. 192.168.1.0/32
- D. 10.0.1.0/24

Answer: D

Explanation:

The allowed block size is between a /28 netmask and /16 netmask. The CIDR block must not overlap with any existing CIDR block that's associated with the VPC.

<https://docs.aws.amazon.com/vpc/latest/userguide/configure-your-vpc.html>

NEW QUESTION 177

- (Topic 3)

A company collects data from a large number of participants who use wearable devices. The company stores the data in an Amazon DynamoDB table and uses applications to analyze the data. The data workload is constant and predictable. The company wants to stay at or below its forecasted budget for DynamoDB. Which solution will meet these requirements MOST cost-effectively?

- A. Use provisioned mode and DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA). Reserve capacity for the forecasted workload.
- B. Use provisioned mode. Specify the read capacity units (RCUs) and write capacity units (WCUs).
- C. Use on-demand mode.
- D. Set the read capacity unit (RCUs) and write capacity units (WCUs) high enough to accommodate changes in the workload.
- E. Use on-demand mode.
- F. Specify the read capacity units (RCUs) and write capacity units (WCUs) with reserved capacity.

Answer: B

Explanation:

This option is the most efficient because it uses provisioned mode, which is a read/write capacity mode for processing reads and writes on your tables that lets you specify how much read and write throughput you expect your application to perform. It also specifies the read capacity units (RCUs) and write capacity units (WCUs), which are the amount of data your application needs to read or write per second. It also meets the requirement of staying at or below its forecasted budget for DynamoDB, as provisioned mode has lower costs than on-demand mode for predictable workloads. This solution meets the requirement of collecting

data from a large number of participants who use wearable devices with a constant and predictable data workload. Option A is less efficient because it uses provisioned mode and DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA), which is a storage class for infrequently accessed items that require milliseconds latency². However, this does not meet the requirement of collecting data from a large number of participants who use wearable devices with a constant and predictable data workload, as DynamoDB Standard-IA is more suitable for items that are accessed less frequently than once every 30 days. Option C is less efficient because it uses on-demand mode, which is a read/write capacity mode that lets you pay only for what you use by automatically adjusting your table's capacity in response to changing demand³. However, this does not meet the requirement of staying at or below its forecasted budget for DynamoDB, as on-demand mode has higher costs than provisioned mode for predictable workloads. Option D is less efficient because it uses on-demand mode and specifies the RCUs and WCUs with reserved capacity, which is a way to reserve read and write capacity for your tables in exchange for discounted hourly rates. However, this does not meet the requirement of staying at or below its forecasted budget for DynamoDB, as on-demand mode has higher costs than provisioned mode for predictable workloads. Also, specifying RCUs and WCUs with reserved capacity is not possible with on-demand mode, as it only applies to provisioned mode.

NEW QUESTION 179

- (Topic 3)

A company needs to provide its employee with secure access to confidential and sensitive files. The company wants to ensure that the files can be accessed only by authorized users. The files must be downloaded securely to the employees' devices.

The files are stored in an on-premises Windows file server. However, due to an increase in remote usage, the file server is out of capacity.

Which solution will meet these requirements?

- A. Migrate the file server to an Amazon EC2 instance in a public subnet
- B. Configure the security group to limit inbound traffic to the employees' IP addresses.
- C. Migrate the files to an Amazon FSx for Windows File Server file system
- D. Integrate the Amazon FSx file system with the on-premises Active Directory. Configure AWS Client VPN.
- E. Migrate the files to Amazon S3, and create a private VPC endpoint
- F. Create a signed URL to allow download.
- G. Migrate the files to Amazon S3, and create a public VPC endpoint. Allow employees to sign on with AWS IAM Identity Center (AWS SSO).

Answer: B

Explanation:

Windows file server is on-premise and we need something to replicate the data to the cloud, the only option we have is AWS FSx for Windows File Server. Also, since the information is confidential and sensitive, we also want to make sure that the appropriate users have access to it in a secure manner.

<https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>

NEW QUESTION 180

- (Topic 3)

A solutions architect is designing the architecture of a new application being deployed to the AWS Cloud. The application will run on Amazon EC2 On-Demand Instances and will automatically scale across multiple Availability Zones. The EC2 instances will scale up and down frequently throughout the day. An Application Load Balancer (ALB) will handle the load distribution. The architecture needs to support distributed session data management. The company is willing to make changes to code if needed.

What should the solutions architect do to ensure that the architecture supports distributed session data management?

- A. Use Amazon ElastiCache to manage and store session data.
- B. Use session affinity (sticky sessions) of the ALB to manage session data.
- C. Use Session Manager from AWS Systems Manager to manage the session.
- D. Use the GetSessionToken API operation in AWS Security Token Service (AWS STS) to manage the session

Answer: A

Explanation:

<https://aws.amazon.com/vi/caching/session-management/>

In order to address scalability and to provide a shared data storage for sessions that can be accessible from any individual web server, you can abstract the HTTP sessions from the web servers themselves. A common solution for this is to leverage an In-Memory Key/Value store such as Redis and Memcached.

ElastiCache offerings for In-Memory key/value stores include ElastiCache for Redis, which can support replication, and ElastiCache for Memcached which does not support replication.

NEW QUESTION 181

- (Topic 3)

An e-commerce company is running a multi-tier application on AWS. The front-end and backend tiers run on Amazon EC2, and the database runs on Amazon RDS for MySQL. The backend tier communicates with the RDS instance. There are frequent calls to return identical database records from the database that are causing performance slowdowns.

Which action should be taken to improve the performance of the backend?

- A. Implement Amazon SNS to store the database calls.
- B. Implement Amazon ElastiCache to cache the large database.
- C. Implement an RDS for MySQL read replica to cache database calls.
- D. Implement Amazon Kinesis Data Firehose to stream the calls to the database.

Answer: B

Explanation:

The best solution is to implement Amazon ElastiCache to cache the large datasets, which will store the frequently accessed data in memory, allowing for faster retrieval times. This can help to alleviate the frequent calls to the database, reduce latency, and improve the overall performance of the backend tier.

NEW QUESTION 185

- (Topic 3)

A company runs a web application that is backed by Amazon RDS. A new database administrator caused data loss by accidentally editing information in a database table. To help recover from this type of incident, the company wants the ability to restore the database to its state from 5 minutes before any change within the last 30 days.

Which feature should the solutions architect include in the design to meet this requirement?

- A. Read replicas
- B. Manual snapshots
- C. Automated backups
- D. Multi-AZ deployments

Answer: C

Explanation:

<https://aws.amazon.com/rds/features/backup/>

Automated backups, will meet the requirement. Amazon RDS allows you to automatically create backups of your DB instance. Automated backups enable point-in-time recovery (PITR) for your DB instance down to a specific second within the retention period, which can be up to 35 days. By setting the retention period to 30 days, the company can restore the database to its state from up to 5 minutes before any change within the last 30 days.

NEW QUESTION 189

- (Topic 3)

A company is building a data analysis platform on AWS by using AWS Lake Formation. The platform will ingest data from different sources such as Amazon S3 and Amazon RDS. The company needs a secure solution to prevent access to portions of the data that contain sensitive information.

- A. Create an IAM role that includes permissions to access Lake Formation tables.
- B. Create data filters to implement row-level security and cell-level security.
- C. Create an AWS Lambda function that removes sensitive information before Lake Formation ingests re data.
- D. Create an AWS Lambda function that periodically Queries and removes sensitive information from Lake Formation tables.

Answer: B

Explanation:

This option is the most efficient because it uses data filters, which are specifications that restrict access to certain data in query results and engines integrated with Lake Formation¹. Data filters can be used to implement row-level security and cell-level security, which are techniques to prevent access to portions of the data that contain sensitive information². Data filters can be applied when granting Lake Formation permissions on a Data Catalog table, and can use PartiQL expressions to filter data based on conditions³. This solution meets the requirement of providing a secure solution to prevent access to portions of the data that contain sensitive information. Option A is less efficient because it uses an IAM role that includes permissions to access Lake Formation tables, which is a way to grant access to data in Lake Formation using IAM policies⁴. However, this does not provide a way to prevent access to portions of the data that contain sensitive information. Option C is less efficient because it uses an AWS Lambda function that removes sensitive information before Lake Formation ingests the data, which is a way to perform data cleansing or transformation using serverless functions. However, this could involve significant changes to the application code and logic, and could also result in data loss or inconsistency. Option D is less efficient because it uses an AWS Lambda function that periodically queries and removes sensitive information from Lake Formation tables, which is a way to perform data cleansing or transformation using serverless functions. However, this could involve significant changes to the application code and logic, and could also result in data loss or inconsistency.

NEW QUESTION 190

- (Topic 3)

A company runs an application on Amazon EC2 Linux instances across multiple Availability Zones. The application needs a storage layer that is highly available and Portable Operating System Interface (POSIX) compliant. The storage layer must provide maximum data durability and must be shareable across the EC2 instances. The data in the storage layer will be accessed frequency for the first 30 days and will be accessed infrequently alter that time.

Which solution will meet these requirements MOST cost-effectively?

- A. Use the Amazon S3 Standard storage class Create an S3 Lifecycle policy to move infrequently accessed data to S3 Glacier
- B. Use the Amazon S3 Standard storage clas
- C. Create an S3 Lifecycle policy to move infrequently accessed data to S3 Standard-Infrequent Access (EF3 Standard-IA).
- D. Use the Amazon Elastic File System (Amazon EFS) Standard storage clas
- E. Create a Lifecycle management policy to move infrequently accessed data to EFS Standard- Infrequent Access (EFS Standard-IA)
- F. Use the Amazon Elastic File System (Amazon EFS) One Zone storage clas
- G. Create a Lifecycle management policy to move infrequently accessed data to EFS One Zone- Infrequent Access (EFS One Zone-IA).

Answer: C

Explanation:

<https://aws.amazon.com/efs/features/infrequent-access/>

NEW QUESTION 194

- (Topic 3)

A company is migrating a Linux-based web server group to AWS. The web servers must access files in a shared file store for some content. The company must not make any changes to the application.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon S3 Standard bucket with access to the web servers.
- B. Configure an Amazon CloudFront distribution with an Amazon S3 bucket as the origin.
- C. Create an Amazon Elastic File System (Amazon EFS) file syste
- D. Mount the EFS file system on all web servers.
- E. Configure a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volum
- F. Mount the EBS volume to all web servers.

Answer: C

Explanation:

Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system on all web servers. To meet the requirements of providing a shared file store for Linux-based web servers without making changes to the application, using an Amazon EFS file system is the best solution. Amazon EFS is a managed NFS file system service that provides shared access to files across multiple Linux-based instances, which makes it suitable for this use case. Amazon S3 is not ideal for this scenario since it is an object storage service and not a file system, and it requires additional tools or libraries to mount the S3 bucket as a file

system. Amazon CloudFront can be used to improve content delivery performance but is not necessary for this requirement. Additionally, Amazon EBS volumes can only be mounted to one instance at a time, so it is not suitable for sharing files across multiple instances.

NEW QUESTION 199

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAA-C03 Practice Exam Features:

- * SAA-C03 Questions and Answers Updated Frequently
- * SAA-C03 Practice Questions Verified by Expert Senior Certified Staff
- * SAA-C03 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * SAA-C03 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAA-C03 Practice Test Here](#)