



Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Topic 1)

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

- A. Create an S3 bucket lifecycle policy to move files from S3 Standard to S3 Glacier 30 days from object creation. Delete the files 4 years after object creation.
- B. Create an S3 bucket lifecycle policy to move files from S3 Standard to S3 One Zone- Infrequent Access (S3 One Zone-IA) 30 days from object creation. Delete the files 4 years after object creation.
- C. Delete the files 4 years after object creation.
- D. Create an S3 bucket lifecycle policy to move files from S3 Standard-Infrequent Access (S3 Standard-IA) 30 days from object creation. Delete the files 4 years after object creation.
- E. Delete the files 4 years after object creation.
- F. Create an S3 bucket lifecycle policy to move files from S3 Standard to S3 Standard- Infrequent Access (S3 Standard-IA) 30 days from object creation. Move the files to S3 Glacier 4 years after object creation.

Answer: B

Explanation:

https://aws.amazon.com/s3/storage-classes/?trk=66264cd8-3b73-416c-9693-ea7cf4fe846a&sc_channel=ps&s_kwid=AL14422!3!536452716950!p!g!aws%20s3%20pricing&ef_id=Cj0KCQjwnbmaBhD-ARIsAGTPcfVHUZN5_BMrzI5zBcaC8KnqpnNZvbZzqPkH6k7q4JcYO5KFLx0YYgaAm6nEALw_wcB:G:s&s_kwid=AL14422!3!536452716950!p!g!aws%20s3%20pricing

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 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 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 4

- (Topic 1)

A company is hosting a static website on Amazon S3 and is using Amazon Route 53 for DNS. The website is experiencing increased demand from around the world. The company must decrease latency for users who access the website.

Which solution meets these requirements MOST cost-effectively?

- A. Replicate the S3 bucket that contains the website to all AWS Regions.
- B. Add Route 53 geolocation routing entries.
- C. Provision accelerators in AWS Global Accelerator.

- D. Associate the supplied IP addresses with the S3 bucket
- E. Edit the Route 53 entries to point to the IP addresses of the accelerators.
- F. Add an Amazon CloudFront distribution in front of the S3 bucket
- G. Edit the Route 53 entries to point to the CloudFront distribution.
- H. Enable S3 Transfer Acceleration on the bucket
- I. Edit the Route 53 entries to point to the new endpoint.

Answer: C

Explanation:

Amazon CloudFront is a content delivery network (CDN) that caches content at edge locations around the world, providing low latency and high transfer speeds to users accessing the content. Adding a CloudFront distribution in front of the S3 bucket will cache the static website's content at edge locations around the world, decreasing latency for users accessing the website. This solution is also cost-effective as it only charges for the data transfer and requests made by users accessing the content from the CloudFront edge locations. Additionally, this solution provides scalability and reliability benefits as CloudFront can automatically scale to handle increased demand and provide high availability for the website.

NEW QUESTION 5

- (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 6

- (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 7

- (Topic 1)

A company needs to keep user transaction data in an Amazon DynamoDB table. The company must retain the data for 7 years. What is the MOST operationally efficient solution that meets these requirements?

- A. Use DynamoDB point-in-time recovery to back up the table continuously.
- B. Use AWS Backup to create backup schedules and retention policies for the table.
- C. Create an on-demand backup of the table by using the DynamoDB console
- D. Store the backup in an Amazon S3 bucket
- E. Set an S3 Lifecycle configuration for the S3 bucket.
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function
- G. Configure the Lambda function to back up the table and to store the backup in an Amazon S3 bucket
- H. Set an S3 Lifecycle configuration for the S3 bucket.

Answer: C

NEW QUESTION 8

- (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 data
- 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 data
- 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 9

- (Topic 1)

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

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

Answer: B

Explanation:

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

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html>

NEW QUESTION 10

- (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 configuration
- 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 configuration
- 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 10

- (Topic 1)

A company wants to improve its ability to clone large amounts of production data into a test environment in the same AWS Region. The data is stored in Amazon EC2 instances on Amazon Elastic Block Store (Amazon EBS) volumes. Modifications to the cloned data must not affect the production environment. The software that accesses this data requires consistently high I/O performance.

A solutions architect needs to minimize the time that is required to clone the production data into the test environment.

Which solution will meet these requirements?

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

Answer: C

Explanation:

To clone the production data into the test environment with high I/O performance and without affecting the production environment, the best option is to take EBS snapshots of the production EBS volumes and restore them onto new EBS volumes in the test environment. Then, attach the new EBS volumes to EC2 instances in the test environment. This option minimizes the time required to clone the data and ensures that modifications to the cloned data do not affect the production environment. Therefore, option C is the correct answer.

Reference: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-restoring-volume.html>

NEW QUESTION 12

- (Topic 1)

A company has an Amazon S3 bucket that contains critical data. The company must protect the data from accidental deletion. Which combination of steps should a solutions architect take to meet these requirements?

(Choose two.)

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

Answer: AB

Explanation:

To protect data in an S3 bucket from accidental deletion, versioning should be enabled, which enables you to preserve, retrieve, and restore every version of every object in an S3 bucket. Additionally, enabling MFA (multi-factor authentication) Delete on the S3 bucket adds an extra layer of protection by requiring an authentication token in addition to the user's access keys to delete objects in the bucket.

Reference:

AWS S3 Versioning documentation: <https://docs.aws.amazon.com/AmazonS3/latest/dev/Versioning.html>

AWS S3 MFA Delete documentation: <https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingMFADelete.html>

NEW QUESTION 16

- (Topic 1)

A company needs to store its accounting records in Amazon S3. The records must be immediately accessible for 1 year and then must be archived for an additional 9 years. No one at the company, including administrative users and root users, can be able to delete the records during the entire 10-year period. The records must be stored with maximum resiliency.

Which solution will meet these requirements?

- A. Store the records in S3 Glacier for the entire 10-year period.
- B. Use an access control policy to deny deletion of the records for a period of 10 years.
- C. Store the records by using S3 Intelligent-Tiering.
- D. Use an IAM policy to deny deletion of the record.
- E. After 10 years, change the IAM policy to allow deletion.
- F. Use an S3 Lifecycle policy to transition the records from S3 Standard to S3 Glacier Deep Archive after 1 year.
- G. Use S3 Object Lock in compliance mode for a period of 10 years.
- H. Use an S3 Lifecycle policy to transition the records from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 1 year.
- I. Use S3 Object Lock in governance mode for a period of 10 years.

Answer: C

Explanation:

To meet the requirements of immediately accessible records for 1 year and then archived for an additional 9 years with maximum resiliency, we can use S3 Lifecycle policy to transition records from S3 Standard to S3 Glacier Deep Archive after 1 year. And to ensure that the records cannot be deleted by anyone, including administrative and root users, we can use S3 Object Lock in compliance mode for a period of 10 years. Therefore, the correct answer is option C.

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html>

NEW QUESTION 18

- (Topic 1)

A bicycle sharing company is developing a multi-tier architecture to track the location of its bicycles during peak operating hours. The company wants to use these data points in its existing analytics platform. A solutions architect must determine the most viable multi-tier option to support this architecture. The data points must be accessible from the REST API.

Which action meets these requirements for storing and retrieving location data?

- A. Use Amazon Athena with Amazon S3
- B. Use Amazon API Gateway with AWS Lambda
- C. Use Amazon QuickSight with Amazon Redshift.
- D. Use Amazon API Gateway with Amazon Kinesis Data Analytics

Answer: D

Explanation:

<https://aws.amazon.com/solutions/implementations/aws-streaming-data-solution-for-amazon-kinesis/>

NEW QUESTION 21

- (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 26

- (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 28

- (Topic 1)

A company uses Amazon S3 to store its confidential audit documents. The S3 bucket uses bucket policies to restrict access to audit team IAM user credentials according to the principle of least privilege. Company managers are worried about accidental deletion of documents in the S3 bucket and want a more secure solution.

What should a solutions architect do to secure the audit documents?

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

Answer: A

NEW QUESTION 31

- (Topic 1)

An application allows users at a company's headquarters to access product data. The product data is stored in an Amazon RDS MySQL DB instance. The operations team has isolated an application performance slowdown and wants to separate read traffic from write traffic. A solutions architect needs to optimize the application's performance quickly.

What should the solutions architect recommend?

- A. Change the existing database to a Multi-AZ deployment
- B. Serve the read requests from the primary Availability Zone.
- C. Change the existing database to a Multi-AZ deployment
- D. Serve the read requests from the secondary Availability Zone.
- E. Create read replicas for the database
- F. Configure the read replicas with half of the compute and storage resources as the source database.
- G. Create read replicas for the database
- H. Configure the read replicas with the same compute and storage resources as the source database.

Answer: D

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.ReadReplicas.html

NEW QUESTION 36

- (Topic 1)

A company wants to reduce the cost of its existing three-tier web architecture. The web, application, and database servers are running on Amazon EC2 instances for the development, test, and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours.

The production EC2 instances run 24 hours a day. The development and test EC2 instances run for at least 8 hours each day. The company plans to implement automation to stop the development and test EC2 instances when they are not in use.

Which EC2 instance purchasing solution will meet the company's requirements MOST cost-effectively?

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

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

Answer: B

NEW QUESTION 41

- (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 VP
- 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 45

- (Topic 2)

A company produces batch data that comes from different databases. The company also produces live stream data from network sensors and application APIs.

The company needs to consolidate all the data into one place for business analytics. The company needs to process the incoming data and then stage the data in different Amazon S3 buckets. Teams will later run one-time queries and import the data into a business intelligence tool to show key performance indicators (KPIs).

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

- A. Use Amazon Athena for one-time queries Use Amazon QuickSight to create dashboards for KPIs
- B. Use Amazon Kinesis Data Analytics for one-time queries Use Amazon QuickSight to create dashboards for KPIs
- C. Create custom AWS Lambda functions to move the individual records from the databases to an Amazon Redshift cluster
- D. Use an AWS Glue extract transform, and load (ETL) job to convert the data into JSON format Load the data into multiple Amazon OpenSearch Service (Amazon Elasticsearch Service) clusters
- E. Use blueprints in AWS Lake Formation to identify the data that can be ingested into a data lake Use AWS Glue to crawl the source extract the data and load the data into Amazon S3 in Apache Parquet format

Answer: AE

Explanation:

Amazon Athena is the best choice for running one-time queries on streaming data. Although Amazon Kinesis Data Analytics provides an easy and familiar standard SQL language to analyze streaming data in real-time, it is designed for continuous queries rather than one-time queries[1]. On the other hand, Amazon Athena is a serverless interactive query service that allows querying data in Amazon S3 using SQL. It is optimized for ad-hoc querying and is ideal for running one-time queries on streaming data[2]. AWS Lake Formation uses as a central place to have all your data for analytics purposes (E). Athena integrates perfectly with S3 and can make queries (A).

NEW QUESTION 46

- (Topic 2)

An application runs on Amazon EC2 instances across multiple Availability Zones. The instances run in an Amazon EC2 Auto Scaling group behind an Application Load Balancer. The application performs best when the CPU utilization of the EC2 instances is at or near 40%.

What should a solutions architect do to maintain the desired performance across all instances in the group?

- A. Use a simple scaling policy to dynamically scale the Auto Scaling group
- B. Use a target tracking policy to dynamically scale the Auto Scaling group
- C. Use an AWS Lambda function to update the desired Auto Scaling group capacity.
- D. Use scheduled scaling actions to scale up and scale down the Auto Scaling group

Answer: B

Explanation:

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

NEW QUESTION 47

- (Topic 2)

An e-commerce company hosts its analytics application in the AWS Cloud. The application generates about 300 MB of data each month. The data is stored in JSON format. The company is evaluating a disaster recovery solution to back up the data. The data must be accessible in milliseconds if it is needed, and the data must be kept for 30 days.

Which solution meets these requirements MOST cost-effectively?

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

Answer: C

Explanation:

This solution meets the requirements of a disaster recovery solution to back up the data that is generated by an analytics application, stored in JSON format, and must be accessible in milliseconds if it is needed. Amazon S3 Standard is a durable and scalable storage class for frequently accessed data. It can store any amount of data and provide high availability and performance. It can also support millisecond access time for data retrieval.

Option A is incorrect because Amazon OpenSearch Service (Amazon Elasticsearch Service) is a search and analytics service that can index and query data, but it is not a backup solution for data stored in JSON format. Option B is incorrect because Amazon S3 Glacier is a low-cost storage class for data archiving and long-term backup, but it does not support millisecond access time for data retrieval. Option D is incorrect because Amazon RDS for PostgreSQL is a relational database service that can store and query structured data, but it is not a backup solution for data stored in JSON format.

References:

? <https://aws.amazon.com/s3/storage-classes/>

? https://aws.amazon.com/s3/faqs/#Durability_and_data_protection

NEW QUESTION 52

- (Topic 2)

A company is migrating its on-premises PostgreSQL database to Amazon Aurora PostgreSQL. The on-premises database must remain online and accessible during the migration. The Aurora database must remain synchronized with the on-premises database.

Which combination of actions must a solutions architect take to meet these requirements? (Choose two.)

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

Answer: AC

Explanation:

AWS Database Migration Service supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle or Microsoft SQL Server to Amazon Aurora. With AWS Database Migration Service, you can also continuously replicate data with low latency from any supported source to any supported target. For example, you can replicate from multiple sources to Amazon Simple Storage Service (Amazon S3) to build a highly available and scalable data lake solution. You can also consolidate databases into a petabyte-scale data warehouse by streaming data to Amazon Redshift. Learn more about the supported source and target databases. <https://aws.amazon.com/dms/>

NEW QUESTION 54

- (Topic 2)

A company runs an application using Amazon ECS. The application creates esi/ed 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 56

- (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 to the 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 bucket
- B. Create an S3 event notification for the analysis S3 bucket
- C. Configure Lambda and SageMaker Pipelines as destinations of the event notification
- D. Configure s3:ObjectCreated:Put as the event type.
- E. Create a Lambda function to copy the files to the analysis S3 bucket
- 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 bucket
- I. Configure Lambda and SageMaker Pipelines as destinations of the event notification
- J. Configure s3:ObjectCreated: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 57

- (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 snapshot
- 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 60

- (Topic 2)

A company is planning to move its data to an Amazon S3 bucket. The data must be encrypted when it is stored in the S3 bucket. Additionally, the encryption key must be automatically rotated every year.

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

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

Answer: B

Explanation:

SSE-S3 - is free and uses AWS owned CMKs (CMK = Customer Master Key). The encryption key is owned and managed by AWS, and is shared among many accounts. Its rotation is automatic with time that varies as shown in the table here. The time is not explicitly defined.

SSE-KMS - has two flavors:

AWS managed CMK. This is free CMK generated only for your account. You can only view its policies and audit usage, but not manage it. Rotation is automatic - once per 1095 days (3 years),

Customer managed CMK. This uses your own key that you create and can manage. Rotation is not enabled by default. But if you enable it, it will be automatically rotated every 1 year. This variant can also use an imported key material by you. If you create such key with an imported material, there is no automated rotation. Only manual rotation.

SSE-C - customer provided key. The encryption key is fully managed by you outside of AWS. AWS will not rotate it.

This solution meets the requirements of moving data to an Amazon S3 bucket, encrypting the data when it is stored in the S3 bucket, and automatically rotating the encryption key every year with the least operational overhead. AWS Key Management Service (AWS KMS) is a service that enables you to create and manage encryption keys for your data. A customer managed key is a symmetric encryption key that you create and manage in AWS KMS. You can enable automatic key rotation for a customer managed key, which means that AWS KMS generates new cryptographic material for the key every year. You can set the S3 bucket's default encryption behavior to use the customer managed KMS key, which means that any object that is uploaded to the bucket without specifying an encryption method will be encrypted with that key.

Option A is incorrect because using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not allow you to control or manage the encryption keys. SSE-S3 uses a unique key for each object, and encrypts that key with a master key that is regularly rotated by S3. However, you cannot enable or disable key rotation for SSE-S3 keys, or specify the rotation interval. Option C is incorrect because manually rotating the KMS key every year can increase the operational overhead and complexity, and it may not meet the requirement of rotating the key every year if you forget or delay the rotation process. Option D is incorrect because encrypting the data with customer key material before moving the data to the S3 bucket can increase the operational overhead and complexity, and it may not provide consistent encryption for all objects in the bucket. Creating a KMS key without key material and importing the customer key material into the KMS key can enable you to use your own source of random bits to generate your KMS keys, but it does not support automatic key

rotation.

References:

? <https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html>

? <https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html>

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-encryption.html>

NEW QUESTION 61

- (Topic 2)

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

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

Answer: A

Explanation:

This solution meets the requirements of a serverless solution, encryption, replication, and SQL analysis with the least operational overhead. Amazon Athena is a serverless interactive query service that can analyze data in S3 using standard SQL. S3 Cross-Region Replication (CRR) can replicate encrypted objects to an S3 bucket in another Region automatically. Server-side encryption with AWS KMS multi-Region keys (SSE-KMS) can encrypt the data at rest using keys that are replicated across multiple Regions. Creating a new S3 bucket can avoid potential conflicts with existing data or configurations.

Option B is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. Option C is incorrect because server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not use KMS keys and it does not support multi-Region replication. Option D is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. It is also incorrect for the same reason as option C. References:

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-4.html>

? <https://aws.amazon.com/blogs/storage/considering-four-different-replication-options-for-data-in-amazon-s3/>

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingEncryption.html>

? <https://aws.amazon.com/athena/>

NEW QUESTION 64

- (Topic 2)

A hospital wants to create digital copies for its large collection of historical written records. The hospital will continue to add hundreds of new documents each day. The hospital's data team will scan the documents and will upload the documents to the AWS Cloud.

A solutions architect must implement a solution to analyze the documents, extract the medical information, and store the documents so that an application can run SQL queries on the data. The solution must maximize scalability and operational efficiency.

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

- A. Write the document information to an Amazon EC2 instance that runs a MySQL database.
- B. Write the document information to an Amazon S3 bucket
- C. Use Amazon Athena to query the data.
- D. Create an Auto Scaling group of Amazon EC2 instances to run a custom application that processes the scanned files and extracts the medical information.
- E. Create an AWS Lambda function that runs when new documents are uploaded
- F. Use Amazon Rekognition to convert the documents to raw text
- G. Use Amazon Transcribe Medical to detect and extract relevant medical information from the text.
- H. Create an AWS Lambda function that runs when new documents are uploaded
- I. Use Amazon Textract to convert the documents to raw text
- J. Use Amazon Comprehend Medical to detect and extract relevant medical information from the text.

Answer: BE

Explanation:

This solution meets the requirements of creating digital copies for a large collection of historical written records, analyzing the documents, extracting the medical information, and storing the documents so that an application can run SQL queries on the data. Writing the document information to an Amazon S3 bucket can provide scalable and durable storage for the scanned files. Using Amazon Athena to query the data can provide serverless and interactive SQL analysis on data stored in S3. Creating an AWS Lambda function that runs when new documents are uploaded can provide event-driven and serverless processing of the scanned files. Using Amazon Textract to convert the documents to raw text can provide accurate optical character recognition (OCR) and extraction of structured data such as tables and forms from documents using artificial intelligence (AI). Using Amazon Comprehend Medical to detect and extract relevant medical information from the text can provide natural language processing (NLP) service that uses machine learning that has been pre-trained to understand and extract health data from medical text.

Option A is incorrect because writing the document information to an Amazon EC2 instance that runs a MySQL database can increase the infrastructure overhead and complexity, and it may not be able to handle large volumes of data. Option C is incorrect because creating an Auto Scaling group of Amazon EC2 instances to run a custom application that processes the scanned files and extracts the medical information can increase the infrastructure overhead and complexity, and it may not be able to leverage existing AI and NLP services such as Textract and Comprehend Medical. Option D is incorrect because using Amazon Rekognition to convert the documents to raw text can provide image and video analysis, but it does not support OCR or extraction of structured data from documents. Using Amazon Transcribe Medical to detect and extract relevant medical information from the text can provide speech-to-text transcription service for medical conversations, but it does not support text analysis or extraction of health data from medical text.

References:

? <https://aws.amazon.com/s3/>

- ? <https://aws.amazon.com/athena/>
- ? <https://aws.amazon.com/lambda/>
- ? <https://aws.amazon.com/textract/>
- ? <https://aws.amazon.com/comprehend/medical/>

NEW QUESTION 68

- (Topic 2)

A solutions architect must design a solution that uses Amazon CloudFront with an Amazon S3 origin to store a static website. The company's security policy requires that all website traffic be inspected by AWS WAF. How should the solutions architect comply with these requirements?

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

Answer: D

Explanation:

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

NEW QUESTION 72

- (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 77

- (Topic 2)

A company has two applications: a sender application that sends messages with payloads to be processed and a processing application intended to receive the messages with payloads. The company wants to implement an AWS service to handle messages between the two applications. The sender application can send about 1,000 messages each hour. The messages may take up to 2 days to be processed. If the messages fail to process, they must be retained so that they do not impact the processing of any remaining messages. Which solution meets these requirements and is the MOST operationally efficient?

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

Answer: C

Explanation:

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

NEW QUESTION 81

- (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 86

- (Topic 2)

A global company is using Amazon API Gateway to design REST APIs for its loyalty club users in the us-east-1 Region and the ap-southeast-2 Region. A solutions architect must design a solution to protect these API Gateway managed REST APIs across multiple accounts from SQL injection and cross-site scripting attacks.

Which solution will meet these requirements with the LEAST amount of administrative effort?

- A. Set up AWS WAF in both Region
- B. Associate Regional web ACLs with an API stage.
- C. Set up AWS Firewall Manager in both Region
- D. Centrally configure AWS WAF rules.
- E. Set up AWS Shield in both Region
- F. Associate Regional web ACLs with an API stage.
- G. Set up AWS Shield in one of the Region
- H. Associate Regional web ACLs with an API stage.

Answer: A

Explanation:

Using AWS WAF has several benefits. Additional protection against web attacks using criteria that you specify. You can define criteria using characteristics of web requests such as the following: Presence of SQL code that is likely to be malicious (known as SQL injection). Presence of a script that is likely to be malicious (known as cross-site scripting). AWS Firewall Manager simplifies your administration and maintenance tasks across multiple accounts and resources for a variety of protections. <https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html>

NEW QUESTION 91

- (Topic 2)

A company runs its ecommerce application on AWS. Every new order is published as a message in a RabbitMQ queue that runs on an Amazon EC2 instance in a single Availability Zone. These messages are processed by a different application that runs on a separate EC2 instance. This application stores the details in a PostgreSQL database on another EC2 instance. All the EC2 instances are in the same Availability Zone.

The company needs to redesign its architecture to provide the highest availability with the least operational overhead.

What should a solutions architect do to meet these requirements?

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

Answer: B

Explanation:

Migrating to Amazon MQ reduces the overhead on the queue management. C and D are dismissed. Deciding between A and B means deciding to go for an AutoScaling group for EC2 or an RDS for Postgress (both multi- AZ). The RDS option has less operational impact, as provide as a service the tools and software required. Consider for instance, the effort to add an additional node like a read replica, to the DB. <https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/active-standby-broker-deployment.html> <https://aws.amazon.com/rds/postgresql/>

NEW QUESTION 96

- (Topic 2)

A solutions architect needs to implement a solution to reduce a company's storage costs. All the company's data is in the Amazon S3 Standard storage class. The company must keep all data for at least 25 years. Data from the most recent 2 years must be highly available and immediately retrievable.

Which solution will meet these requirements?

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

Answer: B

Explanation:

https://aws.amazon.com/about-aws/whats-new/2018/04/announcing-s3-one-zone-infrequent-access-a-new-amazon-s3-storage-class/?nc1=h_ls

NEW QUESTION 101

- (Topic 2)

A company runs a global web application on Amazon EC2 instances behind an Application Load Balancer. The application stores data in Amazon Aurora. The company needs to create a disaster recovery solution and can tolerate up to 30 minutes of downtime and potential data loss. The solution does not need to handle the load when the primary infrastructure is healthy.

What should a solutions architect do to meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 106

- (Topic 2)

A company recently started using Amazon Aurora as the data store for its global ecommerce application. When large reports are run, developers report that the ecommerce application is performing poorly. After reviewing metrics in Amazon CloudWatch, a solutions architect finds that the ReadIOPS and CPU Utilization metrics are spiking when monthly reports run.

What is the MOST cost-effective solution?

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.htm>

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

NEW QUESTION 110

- (Topic 2)

A large media company hosts a web application on AWS. The company wants to start caching confidential media files so that users around the world will have reliable access to the files. The content is stored in Amazon S3 buckets. The company must deliver the content quickly, regardless of where the requests originate geographically.

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 113

- (Topic 2)

A company is running several business applications in three separate VPCs within the us-east-1 Region. The applications must be able to communicate between VPCs. The applications also must be able to consistently send hundreds to gigabytes of data each day to a latency-sensitive application that runs in a single on-premises data center.

A solutions architect needs to design a network connectivity solution that maximizes cost-effectiveness.

Which solution meets those requirements?

- A. Configure three AWS Site-to-Site VPN connections from the data center to AWS. Establish connectivity by configuring one VPN connection for each VPC.
- B. Launch a third-party virtual network appliance in each VPC. Establish an IPsec VPN tunnel between the Data center and each virtual appliance.
- C. Set up three AWS Direct Connect connections from the data center to a Direct Connect gateway in us-east-1. Establish connectivity by configuring each VPC to use one of the Direct Connect connections.
- D. Set up one AWS Direct Connect connection from the data center to AWS.
- E. Create a transit gateway, and attach each VPC to the transit gateway.
- F. Establish connectivity between the Direct Connect connection and the transit gateway.

Answer: D

Explanation:

<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-gateway.html>

NEW QUESTION 115

- (Topic 2)

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

What should the solutions architect do to meet these requirements?

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

Answer: A

Explanation:

Composite alarms determine their states by monitoring the states of other alarms. You can **use composite alarms to reduce alarm noise**. For example, you can create a composite alarm where the underlying metric alarms go into ALARM when they meet specific conditions. You then can set up your composite alarm to go into ALARM and send you notifications when the underlying metric alarms go into ALARM by configuring the underlying metric alarms never to take actions. Currently, composite alarms can take the following actions: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Create_Composite_Alarm.html

NEW QUESTION 117

- (Topic 3)

A company is migrating an old application to AWS. The application runs a batch job every hour and is CPU intensive. The batch job takes 15 minutes on average with an on-premises server. The server has 64 virtual CPU (vCPU) and 512 GiB of memory. Which solution will run the batch job within 15 minutes with the LEAST operational overhead?

- A. Use AWS Lambda with functional scaling.
- B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate.
- C. Use Amazon Lightsail with AWS Auto Scaling.
- D. Use AWS Batch on Amazon EC2.

Answer: D

Explanation:

Use AWS Batch on Amazon EC2. AWS Batch is a fully managed batch processing service that can be used to easily run batch jobs on Amazon EC2 instances. It can scale the number of instances to match the workload, allowing the batch job to be completed in the desired time frame with minimal operational overhead. Using AWS Lambda with Amazon API Gateway - AWS Lambda <https://docs.aws.amazon.com/lambda/latest/dg/services-apigateway.html>
AWS Lambda FAQs <https://aws.amazon.com/lambda/faqs/>

NEW QUESTION 119

- (Topic 3)

A company will deploy a web application on AWS. The company hosts the backend database on Amazon RDS for MySQL with a primary DB instance and five read replicas to support scaling needs. The read replicas must log no more than 1 second behind the primary DB instance. The database routinely runs scheduled stored procedures.

As traffic on the website increases, the replicas experience additional lag during periods of peak load. A solutions architect must reduce the replication lag as much as possible. The solutions architect must minimize changes to the application code and must minimize ongoing overhead.

Which solution will meet these requirements?

Migrate the database to Amazon Aurora MySQL. Replace the read replicas with Aurora Replicas, and configure Aurora Auto Scaling. Replace the stored procedures with Aurora MySQL native functions.

Deploy an Amazon ElastiCache for Redis cluster in front of the database. Modify the application to check the cache before the application queries the database. Replace the stored procedures with AWS Lambda functions.

- A. Migrate the database to a MySQL database that runs on Amazon EC2 instance.
- B. Choose large, compute optimized for all replica nodes.
- C. Maintain the stored procedures on the EC2 instances.
- D. Deploy an Amazon ElastiCache for Redis cluster in front of the database.
- E. Modify the application to check the cache before the application queries the database.
- F. Replace the stored procedures with AWS Lambda functions.
- G. Migrate the database to a MySQL database that runs on Amazon EC2 instance.
- H. Choose large, compute optimized EC2 instances for all replica nodes. Maintain the stored procedures on the EC2 instances.
- I. Migrate the database to Amazon DynamoDB, Provision number of read capacity units (RCUs) to support the required throughput, and configure on-demand capacity scaling.
- J. Replace the stored procedures with DynamoDB streams.

Answer: A

Explanation:

Option A is the most appropriate solution for reducing replication lag without significant changes to the application code and minimizing ongoing operational overhead. Migrating the database to Amazon Aurora MySQL allows for improved replication performance and higher scalability compared to Amazon RDS for MySQL. Aurora Replicas provide faster replication, reducing the replication lag, and Aurora Auto Scaling ensures that there are enough Aurora Replicas to handle the incoming traffic. Additionally, Aurora MySQL native functions can replace the stored procedures, reducing the load on the database and improving performance.

NEW QUESTION 124

- (Topic 3)

A company has hundreds of Amazon EC2 Linux-based instances in the AWS Cloud. Systems administrators have used shared SSH keys to manage the instances. After a recent audit, the company's security team is mandating the removal of all shared keys. A solutions architect must design a solution that provides secure

access to the EC2 instances.

Which solution will meet this requirement with the LEAST amount of administrative overhead?

- A. Use AWS Systems Manager Session Manager to connect to the EC2 instances.
- B. Use AWS Security Token Service (AWS STS) to generate one-time SSH keys on demand.
- C. Allow shared SSH access to a set of bastion instance
- D. Configure all other instances to allow only SSH access from the bastion instances
- E. Use an Amazon Cognito custom authorizer to authenticate user
- F. Invoke an AWS Lambda function to generate a temporary SSH key.

Answer: A

Explanation:

Session Manager is a fully managed AWS Systems Manager capability. With Session Manager, you can manage your Amazon Elastic Compute Cloud (Amazon EC2) instances, edge devices, on-premises servers, and virtual machines (VMs). You can use either an interactive one-click browser-based shell or the AWS Command Line Interface (AWS CLI). Session Manager provides secure and auditable node management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys. Session Manager also allows you to comply with corporate policies that require controlled access to managed nodes, strict security practices, and fully auditable logs with node access details, while providing end users with simple one-click cross-platform access to your managed nodes.
<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

NEW QUESTION 127

- (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 131

- (Topic 3)

A solutions architect is designing the architecture for a software demonstration environment The environment will run on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB) The system will experience significant increases in traffic during working hours but is not required to operate on weekends.

Which combination of actions should the solutions architect take to ensure that the system can scale to meet demand? (Select TWO)

- A. Use AWS Auto Scaling to adjust the ALB capacity based on request rate
- B. Use AWS Auto Scaling to scale the capacity of the VPC internet gateway
- C. Launch the EC2 instances in multiple AWS Regions to distribute the load across Regions
- D. Use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization
- E. Use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends Revert to the default values at the start of the week

Answer: DE

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-target-tracking.html#target-tracking-choose-metrics>

A target tracking scaling policy is a type of dynamic scaling policy that adjusts the capacity of an Auto Scaling group based on a specified metric and a target value¹. A target tracking scaling policy can automatically scale out or scale in the Auto Scaling group to keep the actual metric value at or near the target value¹. A target tracking scaling policy is suitable for scenarios where the load on the application changes frequently and unpredictably, such as during working hours². To meet the requirements of the scenario, the solutions architect should use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization. Instance CPU utilization is a common metric that reflects the demand on the application¹. The solutions architect should specify a target value that represents the ideal average CPU utilization level for the application, such as 50 percent¹. Then, the Auto Scaling group will scale out or scale in to maintain that level of CPU utilization. Scheduled scaling is a type of scaling policy that performs scaling actions based on a date and time³. Scheduled scaling is suitable for scenarios where the load on the application changes periodically and predictably, such as on weekends². To meet the requirements of the scenario, the solutions architect should also use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends. This way, the Auto Scaling group will terminate all instances on weekends when they are not required to operate. The solutions architect should also revert to the default values at the start of the week, so that the Auto Scaling group can resume normal operation.

NEW QUESTION 135

- (Topic 3)

A developer has an application that uses an AWS Lambda function to upload files to Amazon S3 and needs the required permissions to perform the task The developer already has an IAM user with valid IAM credentials required for Amazon S3

What should a solutions architect do to grant the permissions?

- A. Add required IAM permissions in the resource policy of the Lambda function
- B. Create a signed request using the existing IAM credentials in the Lambda function
- C. Create a new IAM user and use the existing IAM credentials in the Lambda function.
- D. Create an IAM execution role with the required permissions and attach the IAM role to the Lambda function

Answer: D

Explanation:

To grant the necessary permissions to an AWS Lambda function to upload files to Amazon S3, a solutions architect should create an IAM execution role with the required permissions and attach the IAM role to the Lambda function. This approach follows the principle of least privilege and ensures that the Lambda function can only access the resources it needs to perform its specific task.

NEW QUESTION 137

- (Topic 3)

An IAM user made several configuration changes to AWS resources in their company's account during a production deployment last week. A solutions architect learned that a couple of security group rules are not configured as desired. The solutions architect wants to confirm which IAM user was responsible for making changes.

Which service should the solutions architect use to find the desired information?

- A. Amazon GuardDuty
- B. Amazon Inspector
- C. AWS CloudTrail
- D. AWS Config

Answer: C

Explanation:

The best option is to use AWS CloudTrail to find the desired information. AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of AWS account activities. CloudTrail can be used to log all changes made to resources in an AWS account, including changes made by IAM users, EC2 instances, AWS management console, and other AWS services. By using CloudTrail, the solutions architect can identify the IAM user who made the configuration changes to the security group rules.

NEW QUESTION 141

- (Topic 3)

A company has a three-tier application on AWS that ingests sensor data from its users' devices. The traffic flows through a Network Load Balancer (NLB) then to Amazon EC2 instances for the web tier and finally to EC2 instances for the application tier. The application tier makes calls to a database. What should a solutions architect do to improve the security of the data in transit?

- A. Configure a TLS listener. Deploy the server certificate on the NLB.
- B. Configure AWS Shield Advanced. Enable AWS WAF on the NLB.
- C. Change the load balancer to an Application Load Balancer (ALB). Enable AWS WAF on the ALB.
- D. Encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances by using AWS Key Management Service (AWS KMS).

Answer: A

Explanation:

The best option to improve the security of the data in transit is to configure a TLS listener and deploy the server certificate on the NLB. This will ensure that the data is encrypted and secure as it travels through the network. Additionally, you could also configure AWS Shield Advanced and enable AWS WAF on the NLB to further protect the network from malicious attacks. Alternatively, you could also change the load balancer to an Application Load Balancer (ALB) and enable AWS WAF on the ALB. Finally, you could also encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances by using AWS Key Management Service (AWS KMS).

You must specify an SSL certificate for a TLS listener. The load balancer uses the certificate to terminate the connection and decrypt requests from clients before routing them to targets. <https://docs.aws.amazon.com/elasticloadbalancing/latest/network/create-listener.html>

NEW QUESTION 142

- (Topic 3)

A company has a regional subscription-based streaming service that runs in a single AWS Region. The architecture consists of web servers and application servers on Amazon EC2 instances. The EC2 instances are in Auto Scaling groups behind Elastic Load Balancers. The architecture includes an Amazon Aurora database cluster that extends across multiple Availability Zones.

The company wants to expand globally and to ensure that its application has minimal downtime.

- A. Extend the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Region.
- B. Use an Aurora global database to deploy the database in the primary Region and the second Region.
- C. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
- D. Deploy the web tier and the application tier to a second Region.
- E. Add an Aurora PostgreSQL cross-Region Aurora Replica in the second Region.
- F. Use Amazon Route 53 health checks with a failovers routing policy to the second Region, Promote the secondary to primary as needed.
- G. Deploy the web tier and the application tier to a second Region.
- H. Create an Aurora PostgreSQL database in the second Region.
- I. Use AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Region.
- J. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
- K. Deploy the web tier and the application tier to a second Region.
- L. Use an Amazon Aurora global database to deploy the database in the primary Region and the second Region.
- M. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
- N. Promote the secondary to primary as needed.

Answer: D

Explanation:

This option is the most efficient because it deploys the web tier and the application tier to a second Region, which provides high availability and redundancy for the

application. It also uses an Amazon Aurora global database, which is a feature that allows a single Aurora database to span multiple AWS Regions¹. It also deploys the database in the primary Region and the second Region, which provides low latency global reads and fast recovery from a Regional outage. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which provides data protection by routing traffic to healthy endpoints in different Regions². It also promotes the secondary to primary as needed, which provides data consistency by allowing write operations in one of the Regions at a time³. This solution meets the requirement of expanding globally and ensuring that its application has minimal downtime. Option A is less efficient because it extends the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Region, which could incur higher costs and complexity than deploying them separately. It also uses an Aurora global database to deploy the database in the primary Region and the second Region, which is correct. However, it does not use Amazon Route 53 health checks with a failover routing policy to the second Region, which could result in traffic being routed to unhealthy endpoints. Option B is less efficient because it deploys the web tier and the application tier to a second Region, which is correct. It also adds an Aurora PostgreSQL cross-Region Aurora Replica in the second Region, which provides read scalability across Regions. However, it does not use an Aurora global database, which provides faster replication and recovery than cross-Region replicas. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which is correct. However, it does not promote the secondary to primary as needed, which could result in data inconsistency or loss. Option C is less efficient because it deploys the web tier and the application tier to a second Region, which is correct. It also creates an Aurora PostgreSQL database in the second Region, which provides data redundancy across Regions. However, it does not use an Aurora global database or cross-Region replicas, which provide faster replication and recovery than creating separate databases. It also uses AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Region, which provides data migration between different sources and targets. However, it does not use an Aurora global database or cross-Region replicas, which provide faster replication and recovery than using AWS DMS. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which is correct.

NEW QUESTION 143

- (Topic 3)

A company is deploying a new application on Amazon EC2 instances. The application writes data to Amazon Elastic Block Store (Amazon EBS) volumes. The company needs to ensure that all data that is written to the EBS volumes is encrypted at rest.

Which solution will meet this requirement?

- A. Create an IAM role that specifies EBS encryption
- B. Attach the role to the EC2 instances.
- C. Create the EBS volumes as encrypted volumes Attach the EBS volumes to the EC2 instances.
- D. Create an EC2 instance tag that has a key of Encrypt and a value of Tru
- E. Tag all instances that require encryption at the ESS level.
- F. Create an AWS Key Management Service (AWS KMS) key policy that enforces EBS encryption in the account Ensure that the key policy is active.

Answer: B

Explanation:

The solution that will meet the requirement of ensuring that all data that is written to the EBS volumes is encrypted at rest is B. Create the EBS volumes as encrypted volumes and attach the encrypted EBS volumes to the EC2 instances. When you create an EBS volume, you can specify whether to encrypt the volume. If you choose to encrypt the volume, all data written to the volume is automatically encrypted at rest using AWS-managed keys. You can also use customer-managed keys (CMKs) stored in AWS KMS to encrypt and protect your EBS volumes. You can create encrypted EBS volumes and attach them to EC2 instances to ensure that all data written to the volumes is encrypted at rest.

NEW QUESTION 148

- (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 152

- (Topic 3)

A solutions architect needs to design a system to store client case files. The files are core company assets and are important. The number of files will grow over time.

The files must be simultaneously accessible from multiple application servers that run on Amazon EC2 instances. The solution must have built-in redundancy.

Which solution meets these requirements?

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon Elastic Block Store (Amazon EBS)
- C. Amazon S3 Glacier Deep Archive
- D. AWS Backup

Answer: A

Explanation:

Amazon EFS provides a simple, scalable, fully managed file system that can be simultaneously accessed from multiple EC2 instances and provides built-in redundancy. It is optimized for multiple EC2 instances to access the same files, and it is designed to be highly available, durable, and secure. It can scale up to petabytes of data and can handle thousands of concurrent connections, and is a cost-effective solution for storing and accessing large amounts of data.

NEW QUESTION 156

- (Topic 3)

A research laboratory needs to process approximately 8 TB of data. The laboratory requires sub-millisecond latencies and a minimum throughput of 6 GBps for the storage subsystem. Hundreds of Amazon EC2 instances that run Amazon Linux will distribute and process the data. Which solution will meet the performance requirements?

- A. Create an Amazon FSx for NetApp ONTAP file system. Set each volume's tiering policy to ALL. Import the raw data into the file system. Mount the file system on the EC2 instances.
- B. Create an Amazon S3 bucket to store the raw data. Create an Amazon FSx for Lustre file system that uses persistent SSD storage. Select the option to import data from and export data to Amazon S3. Mount the file system on the EC2 instances.
- C. Create an Amazon S3 bucket to store the raw data. Create an Amazon FSx for Lustre file system that uses persistent HDD storage. Select the option to import data from and export data to Amazon S3. Mount the file system on the EC2 instances.
- D. Create an Amazon FSx for NetApp ONTAP file system. Set each volume's tiering policy to NON.
- E. Import the raw data into the file system. Mount the file system on the EC2 instances.

Answer: B

Explanation:

Create an Amazon S3 bucket to store the raw data. Create an Amazon FSx for Lustre file system that uses persistent SSD storage. Select the option to import data from and export data to Amazon S3. Mount the file system on the EC2 instances. Amazon FSx for Lustre uses SSD storage for sub-millisecond latencies and up to 6 GBps throughput, and can import data from and export data to Amazon S3. Additionally, the option to select persistent SSD storage will ensure that the data is stored on the disk and not lost if the file system is stopped.

NEW QUESTION 160

- (Topic 3)

A company provides an API to its users that automates inquiries for tax computations based on item prices. The company experiences a larger number of inquiries during the holiday season only that cause slower response times. A solutions architect needs to design a solution that is scalable and elastic. What should the solutions architect do to accomplish this?

- A. Provide an API hosted on an Amazon EC2 instance.
- B. The EC2 instance performs the required computations when the API request is made.
- C. Design a REST API using Amazon API Gateway that accepts the item name.
- D. API Gateway passes item names to AWS Lambda for tax computations.
- E. Create an Application Load Balancer that has two Amazon EC2 instances behind it.
- F. The EC2 instances will compute the tax on the received item names.
- G. Design a REST API using Amazon API Gateway that connects with an API hosted on an Amazon EC2 instance.
- H. API Gateway accepts and passes the item names to the EC2 instance for tax computations.

Answer: B

Explanation:

Lambda server-less is scalable and elastic than EC2 api gateway solution.

NEW QUESTION 165

- (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 166

- (Topic 3)

A company plans to use Amazon ElastiCache for its multi-tier web application. A solutions architect creates a Cache VPC for the ElastiCache cluster and an App VPC for the application's Amazon EC2 instances. Both VPCs are in the us-east-1 Region.

The solutions architect must implement a solution to provide the application's EC2 instances with access to the ElastiCache cluster.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a peering connection between the VPCs. Add a route table entry for the peering connection in both VPCs. Configure an inbound rule for the ElastiCache cluster's security group to allow inbound connection from the application's security group.
- B. Create a Transit VPC. Update the VPC route tables in the Cache VPC and the App VPC to route traffic through the Transit VPC. Configure an inbound rule for the ElastiCache cluster's security group to allow inbound connection from the application's security group.
- C. Create a peering connection between the VPCs. Add a route table entry for the peering connection in both VPCs. Configure an inbound rule for the peering connection's security group to allow inbound connection from the application's security group.
- D. Create a Transit VPC. Update the VPC route tables in the Cache VPC and the App VPC to route traffic through the Transit VPC. Configure an inbound rule for

the Transit VPCs security group to allow inbound connection from the application's security group

Answer: A

Explanation:

Creating a peering connection between the two VPCs and configuring an inbound rule for the ElastiCache cluster's security group to allow inbound connection from the application's security group is the most cost-effective solution. Peering connections are free and you only incur the cost of configuring the security group rules. The Transit VPC solution requires additional VPCs and associated resources, which would incur additional costs.

Before Testing | AWS Certification Information and Policies | AWS <https://aws.amazon.com/certification/policies/before-testing/>

NEW QUESTION 169

- (Topic 3)

A company stores its data objects in Amazon S3 Standard storage. A solutions architect has found that 75% of the data is rarely accessed after 30 days. The company needs all the data to remain immediately accessible with the same high availability and resiliency, but the company wants to minimize storage costs. Which storage solution will meet these requirements?

- A. Move the data objects to S3 Glacier Deep Archive after 30 days.
- B. Move the data objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days.
- C. Move the data objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.
- D. Move the data objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) immediately.

Answer: B

Explanation:

Move the data objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days - will meet the requirements of keeping the data immediately accessible with high availability and resiliency, while minimizing storage costs. S3 Standard-IA is designed for infrequently accessed data, and it provides a lower storage cost than S3 Standard, while still offering the same low latency, high throughput, and high durability as S3 Standard.

NEW QUESTION 171

- (Topic 3)

An ecommerce company is building a distributed application that involves several serverless functions and AWS services to complete order-processing tasks. These tasks require manual approvals as part of the workflow A solutions architect needs to design an architecture for the order-processing application The solution must be able to combine multiple AWS Lambda functions into responsive serverless applications The solution also must orchestrate data and services that run on Amazon EC2 instances, containers, or on- premises servers Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Step Functions to build the application.
- B. Integrate all the application components in an AWS Glue job
- C. Use Amazon Simple Queue Service (Amazon SQS) to build the application
- D. Use AWS Lambda functions and Amazon EventBridge (Amazon CloudWatch Events) events to build the application

Answer: A

Explanation:

AWS Step Functions is a fully managed service that makes it easy to build applications by coordinating the components of distributed applications and microservices using visual workflows. With Step Functions, you can combine multiple AWS Lambda functions into responsive serverless applications and orchestrate data and services that run on Amazon EC2 instances, containers, or on-premises servers. Step Functions also allows for manual approvals as part of the workflow. This solution meets all the requirements with the least operational overhead.

[https://aws.amazon.com/step-functions/#:~:text=AWS%20Step%20Functions%20is%20a,machine%20learning%20\(ML\)%20pipelines.](https://aws.amazon.com/step-functions/#:~:text=AWS%20Step%20Functions%20is%20a,machine%20learning%20(ML)%20pipelines.)

NEW QUESTION 173

- (Topic 3)

A company is developing a new mobile app. The company must implement proper traffic filtering to protect its Application Load Balancer (ALB) against common application-level attacks, such as cross-site scripting or SQL injection. The company has minimal infrastructure and operational staff. The company needs to reduce its share of the responsibility in managing, updating, and securing servers for its AWS environment.

What should a solutions architect recommend to meet these requirements?

- A. Configure AWS WAF rules and associate them with the ALB.
- B. Deploy the application using Amazon S3 with public hosting enabled.
- C. Deploy AWS Shield Advanced and add the ALB as a protected resource.
- D. Create a new ALB that directs traffic to an Amazon EC2 instance running a third-party firewall, which then passes the traffic to the current ALB.

Answer: A

Explanation:

A solutions architect should recommend option A, which is to configure AWS WAF rules and associate them with the ALB. This will allow the company to apply traffic filtering at the application layer, which is necessary for protecting the ALB against common application-level attacks such as cross-site scripting or SQL injection. AWS WAF is a managed service that makes it easy to protect web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. The company can easily manage and update the rules to ensure the security of its application.

NEW QUESTION 178

- (Topic 3)

A company has launched an Amazon RDS for MySQL D6 instance Most of the connections to the database come from serverless applications. Application traffic to the database changes significantly at random intervals At times of high demand, users report that their applications experience database connection rejection errors.

Which solution will resolve this issue with the LEAST operational overhead?

- A. Create a proxy in RDS Proxy Configure the users' applications to use the DB instance through RDS Proxy
- B. Deploy Amazon ElastiCache for Memcached between the users' application and the DB instance

- C. Migrate the DB instance to a different instance class that has higher I/O capacity
- D. Configure the users' applications to use the new DB instance.
- E. Configure Multi-AZ for the DB instance Configure the users' application to switch between the DB instances.

Answer: A

Explanation:

Many applications, including those built on modern serverless architectures, can have a large number of open connections to the database server and may open and close database connections at a high rate, exhausting database memory and compute resources. Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. (<https://aws.amazon.com/pt/rds/proxy/>)

NEW QUESTION 180

- (Topic 3)

A company has a web application that is based on Java and PHP. The company plans to move the application from on-premises to AWS. The company needs the ability to test new site features frequently. The company also needs a highly available and managed solution that requires minimum operational overhead. Which solution will meet these requirements?

- A. Create an Amazon S3 bucket. Enable static web hosting on the S3 bucket. Upload the static content to the S3 bucket. Use AWS Lambda to process all dynamic content.
- B. Deploy the web application to an AWS Elastic Beanstalk environment. Use URL swapping to switch between multiple Elastic Beanstalk environments for feature testing.
- C. Deploy the web application to Amazon EC2 instances that are configured with Java and PHP. Use Auto Scaling groups and an Application Load Balancer to manage the website's availability.
- D. Containerize the web application. Deploy the web application to Amazon EC2 instances. Use the AWS Load Balancer Controller to dynamically route traffic between containers that contain the new site features for testing.

Answer: B

Explanation:

Frequent feature testing -

- Multiple Elastic Beanstalk environments can be created easily for development, testing and production use cases.
- Traffic can be routed between environments for A/B testing and feature iteration using simple URL swapping techniques. No complex routing rules or infrastructure changes required.

NEW QUESTION 184

- (Topic 3)

A company's web application consists of an Amazon API Gateway API in front of an AWS Lambda function and an Amazon DynamoDB database. The Lambda function handles the business logic, and the DynamoDB table hosts the data. The application uses Amazon Cognito user pools to identify the individual users of the application. A solutions architect needs to update the application so that only users who have a subscription can access premium content.

- A. Enable API caching and throttling on the API Gateway API.
- B. Set up AWS WAF on the API Gateway API. Create a rule to filter users who have a subscription.
- C. Apply fine-grained IAM permissions to the premium content in the DynamoDB table.
- D. Implement API usage plans and API keys to limit the access of users who do not have a subscription.

Answer: D

Explanation:

This option is the most efficient because it uses API usage plans and API keys, which are features of Amazon API Gateway that allow you to control who can access your API and how much and how fast they can access it¹. It also implements API usage plans and API keys to limit the access of users who do not have a subscription, which enables you to create different tiers of access for your API and charge users accordingly. This solution meets the requirement of updating the application so that only users who have a subscription can access premium content. Option A is less efficient because it uses API caching and throttling on the API Gateway API, which are features of Amazon API Gateway that allow you to improve the performance and availability of your API and protect your backend systems from traffic spikes². However, this does not provide a way to limit the access of users who do not have a subscription. Option B is less efficient because it uses AWS WAF on the API Gateway API, which is a web application firewall service that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources³. However, this does not provide a way to limit the access of users who do not have a subscription. Option C is less efficient because it uses fine-grained IAM permissions to the premium content in the DynamoDB table, which are permissions that allow you to control access to specific items or attributes within a table⁴. However, this does not provide a way to limit the access of users who do not have a subscription at the API level.

NEW QUESTION 186

- (Topic 3)

A company has implemented a self-managed DNS service on AWS. The solution consists of the following:

- Amazon EC2 instances in different AWS Regions
- Endpoints of a standard accelerator in AWS Global Accelerator

The company wants to protect the solution against DDoS attacks. What should a solutions architect do to meet this requirement?

- A. Subscribe to AWS Shield Advanced. Add the accelerator as a resource to protect.
- B. Subscribe to AWS Shield Advanced. Add the EC2 instances as resources to protect.
- C. Create an AWS WAF web ACL that includes a rate-based rule. Associate the web ACL with the accelerator.
- D. Create an AWS WAF web ACL that includes a rate-based rule. Associate the web ACL with the EC2 instances.

Answer: A

Explanation:

AWS Shield is a managed service that provides protection against Distributed Denial of Service (DDoS) attacks for applications running on AWS. AWS Shield Standard is automatically enabled to all AWS customers at no additional cost. AWS Shield Advanced is an optional paid service. AWS Shield Advanced provides additional protections against more sophisticated and larger attacks for your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, and Route 53. <https://docs.aws.amazon.com/waf/latest/developerguide/ddos-event-mitigation-logic-gax.html>

NEW QUESTION 189

- (Topic 3)

A company is migrating its on-premises workload to the AWS Cloud. The company already uses several Amazon EC2 instances and Amazon RDS DB instances. The company wants a solution that automatically starts and stops the EC2 instances and D6 instances outside of business hours. The solution must minimize cost and infrastructure maintenance.

Which solution will meet these requirements?

- A. Scale the EC2 instances by using elastic resize Scale the DB instances to zero outside of business hours
- B. Explore AWS Marketplace for partner solutions that will automatically start and stop the EC2 Instances and OB instances on a schedule
- C. Launch another EC2 instance
- D. Configure a crontab schedule to run shell scripts that will start and stop the existing EC2 instances and DB instances on a schedule.
- E. Create an AWS Lambda function that will start and stop the EC2 instances and DB instances Configure Amazon EventBridge to invoke the Lambda function on a schedule

Answer: D

Explanation:

The most efficient solution for automatically starting and stopping EC2 instances and DB instances on a schedule while minimizing cost and infrastructure maintenance is to create an AWS Lambda function and configure Amazon EventBridge to invoke the function on a schedule.

Option A, scaling EC2 instances by using elastic resize and scaling DB instances to zero outside of business hours, is not feasible as DB instances cannot be scaled to zero.

Option B, exploring AWS Marketplace for partner solutions, may be an option, but it may not be the most efficient solution and could potentially add additional costs.

Option C, launching another EC2 instance and configuring a crontab schedule to run shell scripts that will start and stop the existing EC2 instances and DB instances on a schedule, adds unnecessary infrastructure and maintenance.

NEW QUESTION 193

- (Topic 3)

A solutions architect wants all new users to have specific complexity requirements and mandatory rotation periods for IAM user passwords. What should the solutions architect do to accomplish this?

- A. Set an overall password policy for the entire AWS account
- B. Set a password policy for each IAM user in the AWS account
- C. Use third-party vendor software to set password requirements
- D. Attach an Amazon CloudWatch rule to the Create_newuser event to set the password with the appropriate requirements

Answer: A

Explanation:

This option is the most efficient because it sets an overall password policy for the entire AWS account, which is a way to specify complexity requirements and mandatory rotation periods for IAM user passwords. It also meets the requirement of setting a password policy for all new users, as the password policy applies to all IAM users in the account. This solution meets the requirement of setting specific complexity requirements and mandatory rotation periods for IAM user passwords. Option B is less efficient because it sets a password policy for each IAM user in the AWS account, which is not possible as password policies can only be set at the account level. Option C is less efficient because it uses third-party vendor software to set password requirements, which is not necessary as IAM provides a built-in way to set password policies. Option D is less efficient because it attaches an Amazon CloudWatch rule to the Create_newuser event to set the password with the appropriate requirements, which is not possible as CloudWatch rules cannot modify IAM user passwords.

NEW QUESTION 197

- (Topic 3)

A company is building a mobile app on AWS. The company wants to expand its reach to millions of users. The company needs to build a platform so that authorized users can watch the company's content on their mobile devices.

What should a solutions architect recommend to meet these requirements?

- A. Publish content to a public Amazon S3 bucket
- B. Use AWS Key Management Service (AWS KMS) keys to stream content.
- C. Set up IPsec VPN between the mobile app and the AWS environment to stream content
- D. Use Amazon CloudFront. Provide signed URLs to stream content.
- E. Set up AWS Client VPN between the mobile app and the AWS environment to stream content.

Answer: C

Explanation:

Amazon CloudFront is a content delivery network (CDN) that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds. CloudFront supports signed URLs that provide authorized access to your content. This feature allows the company to control who can access their content and for how long, providing a secure and scalable solution for millions of users.

NEW QUESTION 200

- (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 the 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 205

- (Topic 3)

A company wants to use Amazon S3 for the secondary copy of its on-premises dataset. The company would rarely need to access this copy. The storage solution's cost should be minimal.

Which storage solution meets these requirements?

- A. S3 Standard
- B. S3 Intelligent-Tiering
- C. S3 Standard-Infrequent Access (S3 Standard-IA)
- D. S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: D

Explanation:

S3 One Zone-IA is a storage class that is designed for data that is accessed less frequently, but requires rapid access when needed. Unlike other S3 Storage Classes which store data in a minimum of three Availability Zones (AZs), S3 One Zone-IA stores data in a single AZ and costs 20% less than S3 Standard-IA. This storage class meets the requirements of the company because it provides a low-cost solution for the secondary copy of its on-premises dataset that would rarely need to be accessed. The other storage classes are either more expensive or not suitable for infrequently accessed data.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html>

NEW QUESTION 210

- (Topic 3)

What should a solutions architect do to ensure that all objects uploaded to an Amazon S3 bucket are encrypted?

- A. Update the bucket policy to deny if the PutObject does not have an s3 x-amz-acl header set
- B. Update the bucket policy to deny if the PutObject does not have an s3:x-amz-aci header set to private.
- C. Update the bucket policy to deny if the PutObject does not have an aws SecureTransport header set to true
- D. Update the bucket policy to deny if the PutObject does not have an x-amz-server-side- encryption header set.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/security/how-to-prevent-uploads-of-unencrypted-objects-to-amazon-s3/#:~:text=Solution%20overview,console%2C%20CLI%2C%20or%20SDK.&text=To%20encrypt%20an%20object%20at,S3%2C%20or%20SSE%2DKMS>.

NEW QUESTION 212

- (Topic 3)

A company runs a containerized application on a Kubernetes cluster in an on-premises data center. The company is using a MongoDB database for data storage. The company wants to migrate some of these environments to AWS, but no code changes or deployment method changes are possible at this time. The company needs a solution that minimizes operational overhead.

Which solution meets these requirements?

- A. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes for compute and MongoDB on EC2 for data storage.
- B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute and Amazon DynamoDB for data storage.
- C. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes for compute and Amazon DynamoDB for data storage.
- D. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate for compute and Amazon DocumentDB (with MongoDB compatibility) for data storage.

Answer: D

Explanation:

Amazon DocumentDB (with MongoDB compatibility) is a fast, reliable, and fully managed database service. Amazon DocumentDB makes it easy to set up, operate, and scale MongoDB-compatible databases in the cloud. With Amazon DocumentDB, you can run the same application code and use the same drivers and tools that you use with MongoDB. <https://docs.aws.amazon.com/documentdb/latest/developerguide/what-is.html>

NEW QUESTION 215

- (Topic 3)

A company needs a backup strategy for its three-tier stateless web application. The web application runs on Amazon EC2 instances in an Auto Scaling group with a dynamic scaling policy that is configured to respond to scaling events. The database tier runs on Amazon RDS for PostgreSQL. The web application does not require temporary local storage on the EC2 instances. The company's recovery point objective (RPO) is 2 hours.

The backup strategy must maximize scalability and optimize resource utilization for this environment.

Which solution will meet these requirements?

- A. Take snapshots of Amazon Elastic Block Store (Amazon EBS) volumes of the EC2 instances and database every 2 hours to meet the RPO.
- B. Configure a snapshot lifecycle policy to take Amazon Elastic Block Store (Amazon EBS) snapshots. Enable automated backups in Amazon RDS to meet the RPO.
- C. Retain the latest Amazon Machine Images (AMIs) of the web and application tiers. Enable automated backups in Amazon RDS and use point-in-time recovery to meet the RPO.

D. Take snapshots of Amazon Elastic Block Store (Amazon EBS) volumes of the EC2 instances every 2 hours Enable automated backups in Amazon RDS and use point-in-time recovery to meet the RPO

Answer: C

Explanation:

Since the application has no local data on instances, AMIs alone can meet the RPO by restoring instances from the most recent AMI backup. When combined with automated RDS backups for the database, this provides a complete backup solution for this environment. The other options involving EBS snapshots would be unnecessary given the stateless nature of the instances. AMIs provide all the backup needed for the app tier. This uses native, automated AWS backup features that require minimal ongoing management: - AMI automated backups provide point-in-time recovery for the stateless app tier. - RDS automated backups provide point-in-time recovery for the database.

NEW QUESTION 217

- (Topic 3)

A company needs to create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster to host a digital media streaming application. The EKS cluster will use a managed node group that is backed by Amazon Elastic Block Store (Amazon EBS) volumes for storage. The company must encrypt all data at rest by using a customer managed key that is stored in AWS Key Management Service (AWS KMS)

Which combination of actions will meet this requirement with the LEAST operational overhead? (Select TWO.)

- A. Use a Kubernetes plugin that uses the customer managed key to perform data encryption.
- B. After creation of the EKS cluster, locate the EBS volume
- C. Enable encryption by using the customer managed key.
- D. Enable EBS encryption by default in the AWS Region where the EKS cluster will be create
- E. Select the customer managed key as the default key.
- F. Create the EKS cluster Create an IAM role that has cowlcity that grants permission to the customer managed ke
- G. Associate the role with the EKS cluster.
- H. Store the customer managed key as a Kubernetes secret in the EKS cluste
- I. Use the customer managed key to encrypt the EBS volumes.

Answer: CD

Explanation:

EBS encryption by default is a feature that enables encryption for all new EBS volumes and snapshots created in a Region1. EBS encryption by default uses a service managed key or a customer managed key that is stored in AWS KMS1. EBS encryption by default is suitable for scenarios where data at rest must be encrypted by using a customer managed key, such as the digital media streaming application in the scenario1.

To meet the requirements of the scenario, the solutions architect should enable EBS encryption by default in the AWS Region where the EKS cluster will be created. The solutions architect should select the customer managed key as the default key for encryption1. This way, all new EBS volumes and snapshots created in that Region will be encrypted by using the customer managed key.

EKS encryption provider support is a feature that enables envelope encryption of Kubernetes secrets in EKS with a customer managed key that is stored in AWS KMS2

. Envelope encryption means that data is encrypted by data encryption keys (DEKs) using AES-GCM; DEKs are encrypted by key encryption keys (KEKs) according to configuration in AWS KMS3. EKS encryption provider support is suitable for scenarios where secrets must be encrypted by using a customer managed key, such as the digital media streaming application in the scenario2.

To meet the requirements of the scenario, the solutions architect should create the EKS cluster and create an IAM role that has a policy that grants permission to the customer managed key. The solutions architect should associate the role with the EKS cluste2r. This way, the EKS cluster can use envelope encryption of Kubernetes secrets with the customer managed key.

NEW QUESTION 219

- (Topic 3)

A company is launching an application on AWS. The application uses an Application Load (ALB) to direct traffic to at least two Amazon EC2 instances in a single target group.

The instances are in an Auto Scaling group for each environment. The company requires a development and a production environment. The production environment will have periods of high traffic.

Which solution will configure the development environment MOST cost-effectively?

- A. Reconfigure the target group in the development environment to have one EC2 instance as a target.
- B. Change the ALB balancing algorithm to least outstanding requests.
- C. Reduce the size of the EC2 instances in both environments.
- D. Reduce the maximum number of EC2 instances in the development environment's Auto Scaling group

Answer: D

Explanation:

This option will configure the development environment in the most cost- effective way as it reduces the number of instances running in the development environment and therefore reduces the cost of running the application. The development environment typically requires less resources than the production environment, and it is unlikely that the development environment will have periods of high traffic that would require a large number of instances. By reducing the maximum number of instances in the development environment's Auto Scaling group, the company can save on costs while still maintaining a functional development environment.

NEW QUESTION 221

- (Topic 3)

A solution architect is designing a company's disaster recovery (DR) architecture. The company has a MySQL database that runs on an Amazon EC2 instance in a private subnet with scheduled backup. The DR design to include multiple AWS Regions.

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

- A. Migrate the MySQL database to multiple EC2 instance
- B. Configure a standby EC2 instance in the DR Region Turn on replication.
- C. Migrate the MySQL database to Amazon RD
- D. Use a Multi-AZ deploymen
- E. Turn on read replication for the primary DB instance in the different Availability Zones.

- F. Migrate the MySQL database to an Amazon Aurora global database
- G. Host the primary DB cluster in the primary Region
- H. Host the secondary DB cluster in the DR Region.
- I. Store the scheduled backup of the MySQL database in an Amazon S3 bucket that is configured for S3 Cross-Region Replication (CRR). Use the data backup to restore the database in the DR Region.

Answer: C

Explanation:

Migrate MySQL database to an Amazon Aurora global database is the best solution because it requires minimal operational overhead. Aurora is a managed service that provides automatic failover, so standby instances do not need to be manually configured. The primary DB cluster can be hosted in the primary Region, and the secondary DB cluster can be hosted in the DR Region. This approach ensures that the data is always available and up-to-date in multiple Regions, without requiring significant manual intervention.

NEW QUESTION 224

.....

Relate Links

100% Pass Your SAA-C03 Exam with ExamBible Prep Materials

<https://www.exambible.com/SAA-C03-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>