

Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



NEW QUESTION 1

A company needs guaranteed Amazon EC2 capacity in three specific Availability Zones in a specific AWS Region for an upcoming event that will last 1 week. What should the company do to guarantee the EC2 capacity?

- A. Purchase Reserved instances that specify the Region needed
- B. Create an On Demand Capacity Reservation that specifies the Region needed
- C. Purchase Reserved instances that specify the Region and three Availability Zones needed
- D. Create an On-Demand Capacity Reservation that specifies the Region and three Availability Zones needed

Answer: D

Explanation:

Explanation

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-capacity-reservations.html>: "When you create a Capacity Reservation, you specify: The Availability Zone in which to reserve the capacity"

NEW QUESTION 2

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

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

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

Answer: AD

NEW QUESTION 3

A company has an application that runs on Amazon EC2 instances and uses an Amazon Aurora database. The EC2 instances connect to the database by using user names and passwords that are stored locally in a file. The company wants to minimize the operational overhead of credential management.

What should a solutions architect do to accomplish this goal?

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

Answer: B

NEW QUESTION 4

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:

Explanation

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

NEW QUESTION 5

An application runs on an Amazon EC2 instance in a VPC. The application processes logs that are stored in an Amazon S3 bucket. The EC2 instance needs to access the S3 bucket without connectivity to the internet.

Which solution will provide private network connectivity to Amazon S3?

- A. Create a gateway VPC endpoint to the S3 bucket.
- B. Stream the logs to Amazon CloudWatch Log
- C. Export the logs to the S3 bucket.
- D. Create an instance profile on Amazon EC2 to allow S3 access.
- E. Create an Amazon API Gateway API with a private link to access the S3 endpoint.

Answer: A

NEW QUESTION 6

A development team runs monthly resource-intensive tests on its general purpose Amazon RDS for MySQL DB instance with Performance Insights enabled. The testing lasts for 48 hours once a month and is the only process that uses the database. The team wants to reduce the cost of running the tests without reducing the compute and memory attributes of the DB instance.

Which solution meets these requirements MOST cost-effectively?

- A. Stop the DB instance when tests are complete
- B. Restart the DB instance when required.
- C. Use an Auto Scaling policy with the DB instance to automatically scale when tests are completed.
- D. Create a snapshot when tests are complete
- E. Terminate the DB instance and restore the snapshot when required.
- F. Modify the DB instance to a low-capacity instance when tests are complete
- G. Modify the DB instance again when required.

Answer: C

NEW QUESTION 7

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

NEW QUESTION 8

A company maintains a searchable repository of items on its website. The data is stored in an Amazon RDS for MySQL database table that contains more than 10 million rows. The database has 2 TB of General Purpose SSD storage. There are millions of updates against this data every day through the company's website. The company has noticed that some insert operations are taking 10 seconds or longer. The company has determined that the database storage performance is the problem.

Which solution addresses this performance issue?

- A. Change the storage type to Provisioned IOPS SSD
- B. Change the DB instance to a memory optimized instance class
- C. Change the DB instance to a burstable performance instance class
- D. Enable Multi-AZ RDS read replicas with MySQL native asynchronous replication.

Answer: A

Explanation:

Explanation

<https://aws.amazon.com/ebs/features/>

"Provisioned IOPS volumes are backed by solid-state drives (SSDs) and are the highest performance EBS volumes designed for your critical, I/O intensive database applications. These volumes are ideal for both IOPS-intensive and throughput-intensive workloads that require extremely low latency."

NEW QUESTION 9

A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet.

Which solutions will meet these requirements? (Choose two.)

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

Answer: AC

NEW QUESTION 10

A company hosts a marketing website in an on-premises data center. The website consists of static documents and runs on a single server. An administrator updates the website content infrequently and uses an SFTP client to upload new documents.

The company decides to host its website on AWS and to use Amazon CloudFront. The company's solutions architect creates a CloudFront distribution. The solutions architect must design the most cost-effective and resilient architecture for website hosting to serve as the CloudFront origin.

Which solution will meet these requirements?

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

Answer: D

NEW QUESTION 10

A solutions architect must design a highly available infrastructure for a website. The website is powered by Windows web servers that run on Amazon EC2 instances. The solutions architect must implement a solution that can mitigate a large-scale DDoS attack that originates from thousands of IP addresses.

Downtime is not acceptable for the website.

Which actions should the solutions architect take to protect the website from such an attack? (Select TWO.)

- A. Use AWS Shield Advanced to stop the DDoS attack.
- B. Configure Amazon GuardDuty to automatically block the attackers.
- C. Configure the website to use Amazon CloudFront for both static and dynamic content.
- D. Use an AWS Lambda function to automatically add attacker IP addresses to VPC network ACLs.
- E. Use EC2 Spot Instances in an Auto Scaling group with a target tracking scaling policy that is set to 80% CPU utilization

Answer: AC

NEW QUESTION 11

A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should be protected throughout the entire application stack, and access to the information should be restricted to certain applications.

Which action should the solutions architect take?

- A. Configure a CloudFront signed URL.
- B. Configure a CloudFront signed cookie.
- C. Configure a CloudFront field-level encryption profile.
- D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

Answer: C

Explanation:

Explanation

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/field-level-encryption.html>

"With Amazon CloudFront, you can enforce secure end-to-end connections to origin servers by using HTTPS. Field-level encryption adds an additional layer of security that lets you protect specific data throughout system processing so that only certain applications can see it."

NEW QUESTION 12

A company wants to manage Amazon Machine Images (AMIs). The company currently copies AMIs to the same AWS Region where the AMIs were created. The company needs to design an application that captures AWS API calls and sends alerts whenever the Amazon EC2 CreateImage API operation is called within the company's account.

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

- A. Create an AWS Lambda function to query AWS CloudTrail logs and to send an alert when a CreateImage API call is detected.
- B. Configure AWS CloudTrail with an Amazon Simple Notification Service (Amazon SNS) notification that occurs when updated logs are sent to Amazon S3. Use Amazon Athena to create a new table and to query on CreateImage when an API call is detected.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for the CreateImage API call. Configure the target as an Amazon Simple Notification Service (Amazon SNS) topic to send an alert when a CreateImage API call is detected.
- D. Configure an Amazon Simple Queue Service (Amazon SQS) FIFO queue as a target for AWS CloudTrail log
- E. Create an AWS Lambda function to send an alert to an Amazon Simple Notification Service (Amazon SNS) topic when a CreateImage API call is detected.

Answer: B

NEW QUESTION 14

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

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

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

Answer: A

Explanation:

Explanation

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

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

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

NEW QUESTION 19

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

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

Answer: A

Explanation:

Explanation

You might want to use Transfer Acceleration on a bucket for various reasons, including the following:

You have customers that upload to a centralized bucket from all over the world.

You transfer gigabytes to terabytes of data on a regular basis across continents.

You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3.

<https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html>

[https://aws.amazon.com/s3/transferacceleration/#:~:text=S3%20Transfer%20Acceleration%20\(S3TA\)%20reduces,to%20S3%20for%20remote%20applications:](https://aws.amazon.com/s3/transferacceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3%20for%20remote%20applications:)

"Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much

as 50-500% for long-distance transfer of larger objects. Customers who have either web or mobile

applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet"

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

"Improved throughput - You can upload parts in parallel to improve throughput."

NEW QUESTION 21

A company needs the ability to analyze the log files of its proprietary application. The logs are stored in JSON format in an Amazon S3 bucket. Queries will be simple and will run on-demand. A solutions architect needs to perform the analysis with minimal changes to the existing architecture. What should the solutions architect do to meet these requirements with the LEAST amount of operational overhead?

- A. Use Amazon Redshift to load all the content into one place and run the SQL queries as needed
- B. Use Amazon CloudWatch Logs to store the logs. Run SQL queries as needed from the Amazon CloudWatch console
- C. Use Amazon Athena directly with Amazon S3 to run the queries as needed
- D. Use AWS Glue to catalog the logs. Use a transient Apache Spark cluster on Amazon EMR to run the SQL queries as needed

Answer: C

Explanation:

Explanation

Amazon Athena can be used to query JSON in S3

NEW QUESTION 24

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

- A. Copy the data so both EBS volumes contain all the documents.
- B. Configure the Application Load Balancer to direct a user to the server with the documents
- C. Copy the data from both EBS volumes to Amazon EFS. Modify the application to save new documents to Amazon EFS

D. Configure the Application Load Balancer to send the request to both servers Return eachdocument from the correct server.

Answer: C

Explanation:

Explanation

Amazon EFS provides file storage in the AWS Cloud. With Amazon EFS, you can create a file system, mount the file system on an Amazon EC2 instance, and then read and write data to and from your file system. You can mount an Amazon EFS file system in your VPC, through the Network File System versions 4.0 and 4.1 (NFSv4) protocol. We recommend using a current generation Linux NFSv4.1 client, such as those found in the latest Amazon Linux, Redhat, and Ubuntu

AMIs, in conjunction with the Amazon EFS Mount Helper. For instructions, see Using the amazon-efsutils Tools.

For a list of Amazon EC2 Linux Amazon Machine Images (AMIs) that support this protocol, see NFS Support. For some AMIs, you'll need to install an NFS client to mount your file system on your Amazon EC2 instance. For instructions, see Installing the NFS Client.

You can access your Amazon EFS file system concurrently from multiple NFS clients, so applications that scale beyond a single connection can access a file system. Amazon EC2 instances running in multiple Availability Zones within the same AWS Region can access the file system, so that many users can access and share a common data source.

NEW QUESTION 29

A company observes an increase in Amazon EC2 costs in its most recent bill

The billing team notices unwanted vertical scaling of instance types for a couple of EC2 instances

A solutions architect needs to create a graph comparing the last 2 months of EC2 costs and perform an in-depth analysis to identify the root cause of the vertical scaling

How should the solutions architect generate the information with the LEAST operational overhead?

- A. Use AWS Budgets to create a budget report and compare EC2 costs based on instance types
- B. Use Cost Explorer's granular filtering feature to perform an in-depth analysis of EC2 costs based on instance types
- C. Use graphs from the AWS Billing and Cost Management dashboard to compare EC2 costs based on instance types for the last 2 months
- D. Use AWS Cost and Usage Reports to create a report and send it to an Amazon S3 bucket Use Amazon QuickSight with Amazon S3 as a source to generate an interactive graph based on instance types.

Answer: B

Explanation:

Explanation

AWS Cost Explorer is a tool that enables you to view and analyze your costs and usage. You can explore your usage and costs using the main graph, the Cost Explorer cost and usage reports, or the Cost Explorer RI reports. You can view data for up to the last 12 months, forecast how much you're likely to spend for the next 12 months, and get recommendations for what Reserved Instances to purchase. You can use Cost Explorer to identify areas that need further inquiry and see trends that you can use to understand your costs. <https://docs.aws.amazon.com/costmanagement/latest/userguide/ce-what-is.html>

NEW QUESTION 30

A development team needs to host a website that will be accessed by other teams. The website contents consist of HTML, CSS, client-side JavaScript, and images Which method is the MOST costeffective for hosting the website?

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

Answer: B

Explanation:

Explanation

In Static Websites, Web pages are returned by the server which are prebuilt.

They use simple languages such as HTML, CSS, or JavaScript.

There is no processing of content on the server (according to the user) in Static Websites. Web pages are returned by the server with no change therefore, static Websites are fast.

There is no interaction with databases.

Also, they are less costly as the host does not need to support server-side processing with different languages.

=====

In Dynamic Websites, Web pages are returned by the server which are processed during runtime means they are not prebuilt web pages but they are built during runtime according to the user's demand.

These use server-side scripting languages such as PHP, Node.js, ASP.NET and many more supported by the server.

So, they are slower than static websites but updates and interaction with databases are possible.

NEW QUESTION 35

A company runs an online marketplace web application on AWS. The application serves hundreds of thousands of users during peak hours. The company needs a scalable, near-real-time solution to share the details of millions of financial transactions with several other internal applications Transactions also need to be processed to remove sensitive data before being stored in a document database for low-latency retrieval.

What should a solutions architect recommend to meet these requirements?

- A. Store the transactions data into Amazon DynamoDB Set up a rule in DynamoDB to remove sensitive data from every transaction upon write Use DynamoDB Streams to share the transactions data with other applications
- B. Stream the transactions data into Amazon Kinesis Data Firehose to store data in Amazon DynamoDB and Amazon S3 Use AWS Lambda integration with Kinesis Data Firehose to remove sensitive data
- C. Other applications can consume the data stored in Amazon S3
- D. Stream the transactions data into Amazon Kinesis Data Streams Use AWS Lambda integration to remove sensitive data from every transaction and then store

the transactions data in Amazon DynamoDB Other applications can consume the transactions data off the Kinesis data stream.

E. Store the batched transactions data in Amazon S3 as file

F. Use AWS Lambda to process every file and remove sensitive data before updating the files in Amazon S3 The Lambda function then stores the data in Amazon DynamoDB Other applications can consume transaction files stored in Amazon S3.

Answer: C

Explanation:

Explanation

The destination of your Kinesis Data Firehose delivery stream. Kinesis Data Firehose can send data records to various destinations, including Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, and any HTTP endpoint that is owned by you or any of your third-party service providers. The following are the supported destinations:

* Amazon OpenSearch Service

* Amazon S3

* Datadog

* Dynatrace

* Honeycomb

* HTTP Endpoint

* Logic Monitor

* MongoDB Cloud

* New Relic

* Splunk

* Sumo Logic

<https://docs.aws.amazon.com/firehose/latest/dev/create-name.html>

<https://aws.amazon.com/kinesis/data-streams/>

Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. KDS can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events.

NEW QUESTION 36

A company is preparing to launch a public-facing web application in the AWS Cloud. The architecture consists of Amazon EC2 instances within a VPC behind an Elastic Load Balancer (ELB). A third-party service is used for the DNS. The company's solutions architect must recommend a solution to detect and protect against large-scale DDoS attacks.

Which solution meets these requirements?

A. Enable Amazon GuardDuty on the account.

B. Enable Amazon Inspector on the EC2 instances.

C. Enable AWS Shield and assign Amazon Route 53 to it.

D. Enable AWS Shield Advanced and assign the ELB to it.

Answer: D

NEW QUESTION 41

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

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

A. Use the EC2 serial console to directly access the terminal interface of each instance for administration.

B. Attach the appropriate IAM role to each existing instance and new instance.

C. Use AWS Systems Manager Session Manager to establish a remote SSH session.

D. Create an administrative SSH key pair.

E. Load the public key into each EC2 instance.

F. Deploy a bastion host in a public subnet to provide a tunnel for administration of each instance.

G. Establish an AWS Site-to-Site VPN connection.

H. Instruct administrators to use their local on-premises machines to connect directly to the instances by using SSH keys across the VPN tunnel.

Answer: B

Explanation:

Explanation

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

NEW QUESTION 43

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

NEW QUESTION 44

A company has thousands of edge devices that collectively generate 1 TB of status alerts each day. Each alert is approximately 2 KB in size. A solutions architect needs to implement a solution to ingest and store the alerts for future analysis. The company wants a highly available solution. However, the company needs to minimize costs and does not want to manage additional infrastructure. Additionally, the company wants to keep 14 days of data available for immediate analysis and archive any data older than 14 days. What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts. Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon S3 bucket. Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days.
- B. Launch Amazon EC2 instances across two Availability Zones and place them behind an Elastic Load Balancer to ingest the alerts. Create a script on the EC2 instances that will store the alerts in an Amazon S3 bucket. Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days.
- C. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts. Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon Elasticsearch Service (Amazon ES) cluster. Set up the Amazon ES cluster to take manual snapshots every day and delete data from the cluster that is older than 14 days.
- D. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to ingest the alerts and set the message retention period to 14 days. Configure consumers to poll the SQS queue, check the age of the message, and analyze the message data as needed. If the message is 14 days old, the consumer should copy the message to an Amazon S3 bucket and delete the message from the SQS queue.

Answer: A

Explanation:

Explanation

<https://aws.amazon.com/kinesis/datafirehose/features/?nc=sn&loc=2#:~:text=into%20Amazon%20S3%2C%20Amazon%20Redshift%2C%20Amazon%20OpenSearch%20Service%2C%20Kinesis,Delivery%20streams>

NEW QUESTION 46

A company has an on-premises application that generates a large amount of time-sensitive data that is backed up to Amazon S3. The application has grown and there are user complaints about internet bandwidth limitations. A solutions architect needs to design a long-term solution that allows for both timely backups to Amazon S3 and with minimal impact on internet connectivity for internal users. Which solution meets these requirements?

- A. Establish AWS VPN connections and proxy all traffic through a VPC gateway endpoint.
- B. Establish a new AWS Direct Connect connection and direct backup traffic through this new connection.
- C. Order daily AWS Snowball devices. Load the data onto the Snowball devices and return the devices to AWS each day.
- D. Submit a support ticket through the AWS Management Console. Request the removal of S3 service limits from the account.

Answer: B

NEW QUESTION 51

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

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

Answer: AB

NEW QUESTION 53

A company has a data ingestion workflow that consists of the following:

An Amazon Simple Notification Service (Amazon SNS) topic for notifications about new data deliveries. An AWS Lambda function to process the data and record metadata. The company observes that the ingestion workflow fails occasionally because of network connectivity issues. When such a failure occurs, the Lambda function does not ingest the corresponding data unless the company manually reruns the job.

Which combination of actions should a solutions architect take to ensure that the Lambda function ingests all data in the future? (Select TWO.)

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

Answer: BE

NEW QUESTION 58

A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size.

Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again.

The company also wants to automate remediation.

What should a solutions architect do to meet these requirements with the LEAST development effort?

- A. Use an Amazon S3 bucket as a secure transfer point.
- B. Use Amazon Inspector to scan the objects in the bucket.
- C. If objects contain PII.
- D. trigger an S3 Lifecycle policy to remove the objects that contain PII.
- E. Use an Amazon S3 bucket as a secure transfer point.

- F. Use Amazon Macie to scan the objects in the bucket
- G. If objects contain PII
- H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- I. Implement custom scanning algorithms in an AWS Lambda function
- J. Trigger the function when objects are loaded into the bucket
- K. If objects contain RII
- L. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- M. Implement custom scanning algorithms in an AWS Lambda function
- N. Trigger the function when objects are loaded into the bucket
- O. If objects contain PII
- P. Use Amazon Simple Email Service (Amazon SES) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects that contain PII.

Answer: B

NEW QUESTION 63

A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that the catalog is stored in a durable location.

What should a solutions architect do to meet these requirements?

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

Answer: A

NEW QUESTION 64

A company is preparing to store confidential data in Amazon S3. For compliance reasons the data must be encrypted at rest. Encryption key usage must be logged for auditing purposes. Keys must be rotated every year.

Which solution meets these requirements and is the MOST operationally efficient?

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

Answer: D

Explanation:

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

When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted.

Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use.

AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by default on customer managed CMKs. When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

NEW QUESTION 67

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

Which design should the solutions architect use?

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

Answer: C

Explanation:

"Create an Amazon SQS queue to hold the jobs that need to be processed. Create an Amazon EC2 Auto Scaling group for the compute application. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue."

In this case we need to find a durable and loosely coupled solution for storing jobs. Amazon SQS is ideal for this use case and can be configured to use dynamic scaling based on the number of jobs waiting in the queue. To configure this scaling you can use the backlog per instance metric with the target value being the acceptable backlog per instance to maintain. You can calculate these numbers as follows: Backlog per instance: To calculate your backlog per instance, start with the `ApproximateNumberOfMessagesQueueAttribute` to determine the length of the SQS queue.

NEW QUESTION 70

A company is running a high performance computing (HPC) workload on AWS across many Linux based Amazon EC2 instances. The company needs a shared

storage system that is capable of sub-millisecond latencies, hundreds of Gbps of throughput and millions of IOPS. Users will store millions of small files. Which solution meets these requirements?

- A. Create an Amazon Elastic File System (Amazon EFS) file system Mount the file system on each of the EC2 instances
- B. Create an Amazon S3 bucket Mount the S3 bucket on each of the EC2 instances
- C. Ensure that the EC2 instances use Amazon Elastic Block Store (Amazon EBS) optimized Mount Provisioned IOPS SSD (io2) EBS volumes with Multi-Attach on each instance
- D. Create an Amazon FSx for Lustre file system
- E. Mount the file system on each of the EC2 instances

Answer: D

NEW QUESTION 71

A company is planning to build a high performance computing (HPC) workload as a service solution that is hosted on AWS. A group of 16 Amazon EC2 Linux instances requires the lowest possible latency for node-to-node communication. The instances also need a shared block device volume for high-performing storage.

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 75

A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage.

The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand.

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 78

A public-facing web application queries a database hosted on an Amazon EC2 instance in a private subnet. A large number of queries involve multiple table joins, and the application performance has been degrading due to an increase in complex queries. The application team will be performing updates to improve performance.

What should a solutions architect recommend to the application team? (Select TWO.)

- A. Cache query data in Amazon SQS
- B. Create a read replica to offload queries
- C. Migrate the database to Amazon Athena
- D. Implement Amazon DynamoDB Accelerator to cache data.
- E. Migrate the database to Amazon RDS

Answer: BE

NEW QUESTION 83

A company is storing sensitive user information in an Amazon S3 bucket. The company wants to provide secure access to this bucket from the application tier running on Amazon EC2 instances inside a VPC.

Which combination of steps should a solutions architect take to accomplish this? (Select TWO.)

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

Answer: BD

NEW QUESTION 84

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? (Select TWO.)

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

Answer: CD

NEW QUESTION 85

An image hosting company uploads its large assets to Amazon S3 Standard buckets. The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again. For the first 30 days after upload, the objects will be accessed frequently. The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent. The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.

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

- A. Move assets to S3 Intelligent-Tiering after 30 days.
- B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.
- C. Configure an S3 Lifecycle policy to clean up expired object delete markers.
- D. Move assets to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days.
- E. Move assets to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

Answer: CD

NEW QUESTION 86

A company has an on-premises MySQL database that handles transactional data. The company is migrating the database to the AWS Cloud. The migrated database must maintain compatibility with the company's applications that use the database. The migrated database also must scale automatically during periods of increased demand.

Which migration solution will meet these requirements?

- A. Use native MySQL tools to migrate the database to Amazon RDS for MySQL. Configure elastic storage scaling.
- B. Migrate the database to Amazon Redshift by using the mysqldump utility. Turn on Auto Scaling for the Amazon Redshift cluster.
- C. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon Aurora. Turn on Aurora Auto Scaling.
- D. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon DynamoDB. Configure an Auto Scaling policy.

Answer: C

NEW QUESTION 88

A solution architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should be protected throughout the entire application stack. End access to the information should be restricted to certain applications.

Which action should the solutions architect take?

- A. Configure a CloudFront signed URL.
- B. Configure a CloudFront signed cookie.
- C. Configure a CloudFront field-level encryption profile.
- D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

Answer: C

NEW QUESTION 89

A company has migrated a two-tier application from its on-premises data center to the AWS Cloud. The data tier is a Multi-AZ deployment of Amazon RDS for Oracle with 12 TB of General Purpose SSD Amazon Elastic Block Store (Amazon EBS) storage. The application is designed to process and store documents in the database as binary large objects (blobs) with an average document size of 6 MB.

The database size has grown over time, reducing the performance and increasing the cost of storage. The company must improve the database performance and needs a solution that is highly available and resilient.

Which solution will meet these requirements MOST cost-effectively?

- A. Reduce the RDS DB instance size. Increase the storage capacity to 24 TiB. Change the storage type to Magnetic.
- B. Increase the RDS DB instance size.
- C. Increase the storage capacity to 24 TiB. Change the storage type to Provisioned IOPS.
- D. Create an Amazon S3 bucket.
- E. Update the application to store documents in the S3 bucket. Store the object metadata in the existing database.
- F. Create an Amazon DynamoDB table.
- G. Update the application to use DynamoDB.
- H. Use AWS Database Migration Service (AWS DMS) to migrate data from the Oracle database to DynamoDB.

Answer: C

NEW QUESTION 93

A company uses Amazon EC2 instances to host its internal systems. As part of a deployment operation, an administrator tries to use the AWS CLI to terminate an EC2 instance. However, the administrator receives a 403 (Access Denied) error message. The administrator is using an IAM role that has the following IAM policy attached:

What is the cause of the unsuccessful request?

- A. The EC2 Instance has a resource-based policy with a Deny statement.
- B. The principal has not been specified in the policy statement
- C. The "Action" field does not grant the actions that are required to terminate the EC2 instance
- C. The request to terminate the EC2 instance does not originate from the CIDR blocks 192.0.2.0/24 or 203.0.113.0/24.

Answer: B

NEW QUESTION 94

A company's website handles millions of requests each day and the number of requests continues to increase. A solutions architect needs to improve the response time of the web application. The solutions architect determines that the application needs to decrease latency when retrieving product details from the Amazon DynamoDB table

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

- A. Set up a DynamoDB Accelerator (DAX) cluster Route all read requests through DAX.
- B. Set up Amazon ElastiCache for Redis between the DynamoDB table and the web application Route all read requests through Redis.
- C. Set up Amazon ElastiCache for Memcached between the DynamoDB table and the web application Route all read requests through Memcached.
- D. Set up Amazon DynamoDB streams on the table and have AWS Lambda read from the table and populate Amazon ElastiCache Route all read requests through ElastiCache

Answer: A

NEW QUESTION 97

A company is running a critical business application on Amazon EC2 instances behind an Application Load Balancer The EC2 instances run in an Auto Scaling group and access an Amazon RDS DB instance

The design did not pass an operational review because the EC2 instances and the DB instance are all located in a single Availability Zone A solutions architect must update the design to use a second Availability Zone

Which solution will make the application highly available?

- A. Provision a subnet in each Availability Zone Configure the Auto Scaling group to distribute the EC2 instances across both Availability Zones Configure the DB instance with connections to each network
- B. Provision two subnets that extend across both Availability Zones Configure the Auto Scaling group to distribute the EC2 instances across both Availability Zones Configure the DB instance with connections to each network
- C. Provision a subnet in each Availability Zone Configure the Auto Scaling group to distribute the EC2 instances across both Availability Zones Configure the DB instance for Multi-AZ deployment
- D. Provision a subnet that extends across both Availability Zones Configure the Auto Scaling group to distribute the EC2 instances across both Availability Zones Configure the DB instance for Multi-AZ deployment

Answer: C

NEW QUESTION 98

A company's order system sends requests from clients to Amazon EC2 instances The EC2 instances process the orders and then store the orders in a database on Amazon RDS. Users report that they must reprocess orders when the system fails. The company wants a resilient solution that can process orders automatically if a system outage occurs.

What should a solutions architect do to meet these requirements?

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

Answer: C

NEW QUESTION 103

A company is building a containerized application on premises and decides to move the application to AWS. The application will have thousands of users soon after it is deployed. The company is unsure how to manage the deployment of containers at scale. The company needs to deploy the containerized application in a highly available architecture that minimizes operational overhead.

Which solution will meet these requirements?

- A. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository
- B. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type to run the container
- C. Use target tracking to scale automatically based on demand.
- D. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository
- E. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type to run the container
- F. Use target tracking to scale automatically based on demand.
- G. Store container images in a repository that runs on an Amazon EC2 instance
- H. Run the containers on EC2 instances that are spread across multiple Availability Zones
- I. Monitor the average CPU utilization in Amazon CloudWatch
- J. Launch new EC2 instances as needed
- K. Create an Amazon EC2 Amazon Machine Image (AMI) that contains the container image Launch EC2 instances in an Auto Scaling group across multiple Availability Zones
- L. Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

Answer: A

NEW QUESTION 107

A company needs to store data in Amazon S3 and must prevent the data from being changed. The company wants new objects that are uploaded to Amazon S3 to remain unchangeable for a nonspecific amount of time until the company decides to modify the objects. Only specific users in the company's AWS account can have the ability to delete the objects. What should a solutions architect do to meet these requirements?

- A. Create an S3 Glacier vault Apply a write-once, read-many (WORM) vault lock policy to the objects
- B. Create an S3 bucket with S3 Object Lock enabled Enable versioning Set a retention period of 100 years Use governance mode as the S3 bucket's default retention mode for new objects
- C. Create an S3 bucket Use AWS CloudTrail to track any S3 API events that modify the objects Upon notification, restore the modified objects from any backup versions that the company has
- D. Create an S3 bucket with S3 Object Lock enabled Enable versioning Add a legal hold to the objects Add the s3 PutObjectLegalHold permission to the IAM policies of users who need to delete the objects

Answer: D

NEW QUESTION 109

A company has a web-based map application that provides status information about ongoing repairs. The application sometimes has millions of users. Repair teams have a mobile app that sends current location and status in a JSON message to a REST-based endpoint. Few repairs occur on most days. The company wants the application to be highly available and to scale when large numbers of repairs occur after nature disasters. Customer use the application most often during these times. The company does not want to pay for idle capacity.

- A. Create a webpage that is based on Amazon S3 to display information
- B. Use Amazon API Gateway and AWS Lambda to receive the JSON status data Store the JSON data in Amazon S3.
- C. Use Amazon EC2 instances as web servers across multiple Availability Zones
- D. Run the EC2 instances in an Auto Scaling group
- E. Use Amazon API Gateway and AWS Lambda to receive the JSON status data Store the JSON data in Amazon S3.
- F. Use Amazon EC2 instances as web servers across multiple Availability Zones
- G. Run the EC2 instances in an Auto Scaling group
- H. Use a REST endpoint on the EC2 instances to receive the JSON status data
- I. Store the JSON data in an Amazon RDS Multi-AZ DB instance.
- J. Use Amazon EC2 instances as web servers across multiple Availability Zones Run the EC2 instances in an Auto Scaling group Use a REST endpoint on the EC2 instances to receive the JSON status data Store the JSON data in an Amazon DynamoDB table.

Answer: D

NEW QUESTION 111

A company has a web application that is based on Java and PHP. The company wants 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: D

NEW QUESTION 112

A company is building an e-commerce application and needs to store sensitive customer information. The company needs to give customers the ability to complete purchase transactions on the website. The company also needs to ensure that sensitive customer data is protected, even from database administrators. Which solution meets these requirements?

- A. Store sensitive data in an Amazon Elastic Block Store (Amazon EBS) volume
- B. Use EBS encryption to encrypt the data
- C. Use an IAM instance role to restrict access.
- D. Store sensitive data in Amazon RDS for MySQL
- E. Use AWS Key Management Service (AWS KMS) client-side encryption to encrypt the data.
- F. Store sensitive data in Amazon S3. Use AWS Key Management Service (AWS KMS) service-side encryption of the data
- G. Use S3 bucket policies to restrict access.
- H. Store sensitive data in Amazon FSx for Windows Server
- I. Mount the file share on application servers. Use Windows file permissions to restrict access.

Answer: C

NEW QUESTION 117

A company is experiencing sudden increases in demand. The company needs to provision large Amazon EC2 instances from an Amazon Machine Image (AMI). The instances will run in an Auto Scaling group. The company needs a solution that provides minimum initialization latency to meet the demand. Which solution meets these requirements?

- A. Use the `aws ec2 register-image` command to create an AMI from a snapshot Use AWS Step Functions to replace the AMI in the Auto Scaling group
- B. Enable Amazon Elastic Block Store (Amazon EBS) fast snapshot restore on a snapshot Provision an AMI by using the snapshot Replace the AMI in the Auto Scaling group with the new AMI
- C. Enable AMI creation and define lifecycle rules in Amazon Data Lifecycle Manager (Amazon DLM) Create an AWS Lambda function that modifies the AMI in the Auto Scaling group
- D. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke AWS Backup lifecycle policies that provision AMIs Configure Auto Scaling group capacity

limits as an event source in EventBridge (CloudWatch Events)

Answer: B

NEW QUESTION 121

A company runs a high performance computing (HPC) workload on AWS. The workload required low-latency network performance and high network throughput with tightly coupled node-to-node communication. The Amazon EC2 instances are properly sized for compute and storage capacity, and are launched using default options.

What should a solutions architect propose to improve the performance of the workload?

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

Answer: A

Explanation:

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

NEW QUESTION 123

A gaming company hosts a browser-based application on AWS. The users of the application consume a large number of videos and images that are stored in Amazon S3. This content is the same for all users.

The application has increased in popularity, and millions of users worldwide are accessing these media files. The company wants to provide the files to the users while reducing the load on the origin.

Which solution meets these requirements MOST cost-effectively?

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

Answer: B

Explanation:

CloudFront uses Edge Locations to cache content while Global Accelerator uses Edge Locations to find an optimal pathway to the nearest regional endpoint.

NEW QUESTION 125

A company uses Amazon S3 as its data lake. The company has a new partner that must use SFTP to upload data files. A solutions architect needs to implement a highly available SFTP solution that minimizes operational overhead.

Which solution will meet these requirements?

- A. Use AWS Transfer Family to configure an SFTP-enabled server with a publicly accessible endpoint. Choose the S3 data lake as the destination.
- B. Use Amazon S3 File Gateway as an SFTP server. Expose the S3 File Gateway endpoint URL to the new partner. Share the S3 File Gateway endpoint with the new partner.
- C. Launch an Amazon EC2 instance in a private subnet in a VPC. Instruct the new partner to upload files to the EC2 instance by using a VPN. Run a cron job script on the EC2 instance to upload files to the S3 data lake.
- D. Launch Amazon EC2 instances in a private subnet in a VPC. Place a Network Load Balancer (NLB) in front of the EC2 instances. Create an SFTP listener port for the NLB. Share the NLB hostname with the new partner.
- E. Run a cron job script on the EC2 instances to upload files to the S3 data lake.

Answer: A

NEW QUESTION 127

Availability Zone The company wants the application to be highly available with minimum downtime and minimum loss of data.

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

- A. Place the EC2 instances in different AWS Regions. Use Amazon Route 53 health checks to redirect traffic. Use Aurora PostgreSQL Cross-Region Replication.
- B. Configure the Auto Scaling group to use multiple Availability Zones. Configure the database as Multi-AZ. Configure an Amazon RDS Proxy instance for the database.
- C. Configure the Auto Scaling group to use one Availability Zone. Generate hourly snapshots of the database. Recover the database from the snapshots in the event of a failure.
- D. Configure the Auto Scaling group to use multiple AWS Regions. Write the data from the application to Amazon S3. Use S3 Event Notifications to launch an AWS Lambda function to write the data to the database.

Answer: B

NEW QUESTION 130

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

Which combination of changes will meet those requirements with the LEAST operational overhead? (Select TWO)

- A. Deploy an AWS WAF web ACL in front of the website to provide HTTPS functionality.
- B. Create and deploy an AWS Lambda function to manage and serve the website content.
- C. Create the new website and an Amazon S3 bucket. Deploy the website on the S3 bucket with static website hosting enabled.
- D. Create the new website.

E. Deploy the website by using an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer.

Answer: D

NEW QUESTION 132

A company is designing a new web application that the company will deploy into a single AWS Region. The application requires a two-tier architecture that will include Amazon EC2 instances and an Amazon RDS DB instance. A solutions architect needs to design the application so that all components are highly available.

- A. Deploy EC2 instances in an additional Region Create a DB instance with the Multi-AZ option activated
- B. Deploy all EC2 instances in the same Region and the same Availability Zone
- C. Create a DB instance with the Multi-AZ option activated.
- D. Deploy the EC2 instances across at least two Availability Zones within the same Region
- E. Create a DB instance in a single Availability Zone
- F. Deploy the EC2 instances across at least Two Availability Zones within the same Region
- G. Create a DB instance with the Multi-AZ option activated

Answer: D

NEW QUESTION 136

A company is hosting a website from an Amazon S3 bucket that is configured for public hosting. The company's security team mandates the usage of secure connections for access to the website. However; HTTP-based URLs and HTTPS-based URLs must be functional. What should a solution architect recommend to meet these requirements?

- A. Create an S3 bucket policy to explicitly deny non-HTTPS traffic.
- B. Enable S3 Transfer Acceleration
- C. Select the HTTPS Only bucket property.
- D. Place the website behind an Elastic Load Balancer that is configured to redirect HTTP traffic to HTTPS.
- E. Serve the website through an Amazon CloudFront distribution that is configured to redirect HTTP traffic to HTTPS.

Answer: D

NEW QUESTION 138

A company has a business system that generates hundreds of reports each day. The business system saves the reports to a network share in CSV format. The company needs to store this data in the AWS Cloud in near-real time for analysis. Which solution will meet these requirements with the LEAST administrative overhead?

- A. Use AWS DataSync to transfer the files to Amazon S3 Create a scheduled task that runs at the end of each day.
- B. Create an Amazon S3 File Gateway Update the business system to use a new network share from the S3 File Gateway.
- C. Use AWS DataSync to transfer the files to Amazon S3 Create an application that uses the DataSync API in the automation workflow.
- D. Deploy an AWS Transfer for SFTP endpoint Create a script that checks for new files on the network share and uploads the new files by using SFTP.

Answer: B

NEW QUESTION 139

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: C

NEW QUESTION 142

A company is developing a new machine learning (ML) model solution on AWS. The models are developed as independent microservices that fetch approximately 1GB of model data from Amazon S3 at startup and load the data into memory. Users access the models through an asynchronous API. Users can send a request or a batch of requests and specify where the results should be sent.

The company provides models to hundreds of users. The usage patterns for the models are irregular. Some models could be unused for days or weeks. Other models could receive batches of thousands of requests at a time.

Which design should a solutions architect recommend to meet these requirements?

- A. Direct the requests from the API to a Network Load Balancer (NLB). Deploy the models as AWS Lambda functions that are invoked by the NLB.
- B. Direct the requests from the API to an Application Load Balancer (ALB). Deploy the models as Amazon Elastic Container Service (Amazon ECS) services that read from an Amazon Simple Queue Service (Amazon SQS) queue. Use AWS App Mesh to scale the instances of the ECS cluster based on the SQS queue size.
- C. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue. Deploy the models as AWS Lambda functions that are invoked by SQS events. Use AWS Auto Scaling to increase the number of vCPUs for the Lambda functions based on the SQS queue size.
- D. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue. Deploy the models as Amazon Elastic Container Service (Amazon ECS) services that read from the queue. Enable AWS Auto Scaling on Amazon ECS for both the cluster and copies of the service based on the queue size.

Answer: C

NEW QUESTION 146

A company hosts an application on AWS. The application uses AWS Lambda functions and stores data in Amazon DynamoDB tables. The Lambda functions are connected to a VPC that does not have internet access.

The traffic to access DynamoDB must not travel across the internet. The application must have write access to only specific DynamoDB tables.

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

- A. Attach a VPC endpoint policy for DynamoDB to allow write access to only the specific DynamoDB tables.
- B. Attach a security group to the interface VPC endpoint to allow write access to only the specific DynamoDB tables.
- C. Create a resource-based IAM policy to grant write access to only the specific DynamoDB table
- D. Attach the policy to the DynamoDB tables.
- E. Create a gateway VPC endpoint for DynamoDB that is associated with the Lambda VPC
- F. Ensure that the Lambda execution role can access the gateway VPC endpoint.
- G. Create an interface VPC endpoint for DynamoDB that is associated with the Lambda VPC
- H. Ensure that the Lambda execution role can access the interface VPC endpoint.

Answer: AD

NEW QUESTION 151

A company has developed a new content-sharing application that runs on Amazon Elastic Container Service (Amazon ECS). The application runs on Amazon Linux Docker tasks that use the Amazon EC2 launch type. The application requires a storage solution that has the following characteristics:

- Accessibility (or multiple ECS tasks through bind mounts)
- Resiliency across Availability Zones
- Burstable throughput of up to 3 Gbps
- Ability to be scaled up over time

Which storage solution meets these requirements?

- A. Launch an Amazon FSx for Windows File Server Multi-AZ instance
- B. Configure the ECS task definitions to mount the Amazon FSx instance volume at launch.
- C. Launch an Amazon Elastic File System (Amazon EFS) instance
- D. Configure the ECS task definitions to mount the EFS Instance volume at launch.
- E. Create a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach set to enable
- F. Attach the EBS volume to the ECS EC2 instance Configure ECS task definitions to mount the EBS instance volume at launch.
- G. Launch an EC2 instance with several Provisioned IOPS SSD (k>2) Amazon Elastic Block Store (Amazon EBS) volumes attached in a RAID 0 configuration
- H. Configure the EC2 instance as an NFS storage server
- I. Configure ECS task definitions to mount the volumes at launch.

Answer: B

NEW QUESTION 153

A company needs to move data from an Amazon EC2 instance to an Amazon S3 bucket. The company must ensure that no API calls and no data are routed through public internet routes. Only the EC2 instance can have access to upload data to the S3 bucket.

Which solution will meet these requirements?

- A. Create an interface VPC endpoint for Amazon S3 in the subnet where the EC2 instance is located. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
- B. Create a gateway VPC endpoint for Amazon S3 in the Availability Zone where the EC2 instance is located. Attach appropriate security groups to the endpoint. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
- C. Run the nslookup tool from inside the EC2 instance to obtain the private IP address of the S3 bucket's service API endpoint. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucket. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
- D. Use the AWS provided publicly available IP ranges JSON file to obtain the private IP address of the S3 bucket's service API endpoint. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucket. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.

Answer: B

NEW QUESTION 156

A company wants to build a data lake on AWS from data that is stored in an on-premises Oracle relational database. The data lake must receive ongoing updates from the on-premises database.

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

- A. Use AWS DataSync to transfer the data to Amazon S3. Use AWS Glue to transform the data and integrate the data into a data lake.
- B. Use AWS Snowball to transfer the data to Amazon S3. Use AWS Batch to transform the data and integrate the data into a data lake.
- C. Use AWS Database Migration Service (AWS DMS) to transfer the data to Amazon S3. Use AWS Glue to transform the data and integrate the data into a data lake.
- D. Use an Amazon EC2 instance to transfer the data to Amazon S3. Configure the EC2 instance to transform the data and integrate the data into a data lake.

Answer: C

NEW QUESTION 161

A company wants to migrate its existing on-premises monolithic application to AWS.

The company wants to keep as much of the front-end code and the backend code as possible. However, the company wants to break the application into smaller applications. A different team will manage each application. The company needs a highly scalable solution that minimizes operational overhead.

Which solution will meet these requirements?

- A. Host the application on AWS Lambda. Integrate the application with Amazon API Gateway.
- B. Host the application with AWS Amplify.
- C. Connect the application to an Amazon API Gateway API that is integrated with AWS Lambda.
- D. Host the application on Amazon EC2 instance.

- E. Set up an Application Load Balancer with EC2 instances in an Auto Scaling group as targets.
- F. Host the application on Amazon Elastic Container Service (Amazon ECS) Set up an Application Load Balancer with Amazon ECS as the target.

Answer: C

NEW QUESTION 165

A gaming company wants to launch a new internet-facing application in multiple AWS Regions. The application will use the TCP and UDP protocols for communication. The company needs to provide high availability and minimum latency for global users. Which combination of actions should a solutions architect take to meet these requirements? (Select TWO.)

- A. Create internal Network Load Balancers in front of the application in each Region
- B. Create external Application Load Balancers in front of the application in each Region
- C. Create an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region
- D. Configure Amazon Route 53 to use a geolocation routing policy to distribute the traffic
- E. Configure Amazon CloudFront to handle the traffic and route requests to the application in each Region

Answer: AC

NEW QUESTION 169

An online photo application lets users upload photos and perform image editing operations. The application offers two classes of service: free and paid. Photos submitted by paid users are processed before those submitted by free users. Photos are uploaded to Amazon S3 and the job information is sent to Amazon SQS. Which configuration should a solutions architect recommend?

- A. Use one SQS FIFO queue. Assign a higher priority to the paid photos so they are processed first.
- B. Use two SQS FIFO queues: one for paid and one for free. Set the free queue to use short polling and the paid queue to use long polling.
- C. Use two SQS standard queues: one for paid and one for free. Configure Amazon EC2 instances to prioritize polling for the paid queue over the free queue.
- D. Use one SQS standard queue.
- E. Set the visibility timeout of the paid photos to zero. Configure Amazon EC2 instances to prioritize visibility settings so paid photos are processed first.

Answer: C

Explanation:

<https://acloud.guru/forums/guru-of-the-week/discussion/-L7Be8rOao3InQxdQcXj/> <https://aws.amazon.com/sqs/features/priority/>
Priority: Use separate queues to provide prioritization of work. <https://aws.amazon.com/sqs/features/priority/>
<https://aws.amazon.com/sqs/features/#:~:text=Priority%3A%20Use%20separate%20queues%20to%20provide%20visibility,timeout%3A%20Set%20the%20visibility%20timeout%20of%20the%20paid%20photos%20to%20zero,Configure%20Amazon%20EC2%20instances%20to%20prioritize%20visibility%20settings%20so%20paid%20photos%20are%20processed%20first.>
[https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-short-and-long-polling.](https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-short-and-long-polling.html)

NEW QUESTION 172

A company wants to use Amazon S3 for the secondary copy of its 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: C

NEW QUESTION 175

A company runs a two-tier e-commerce website on AWS. The web tier consists of a load balancer that sends traffic to Amazon EC2 instances. The database tier uses an Amazon RDS D8 instance. The EC2 instances and the RDS DB instance should not be exposed to the public internet. The EC2 instances require internet access to complete payment processing of orders through a third-party web service. The application must be highly available. Which combination of configuration options will meet these requirements? (Select TWO.)

- A. Use an Auto Scaling group to launch the EC2 instances in private subnets. Deploy an RDS Multi-AZ DB instance in private subnets.
- B. Configure a VPC with two private subnets and two NAT gateways across two Availability Zones. Deploy an Application Load Balancer in the private subnets.
- C. Use an Auto Scaling group to launch the EC2 instances in public subnets across two Availability Zones. Deploy an RDS Multi-AZ DB instance in private subnets.
- D. Configure a VPC with one public subnet, one private subnet, and two NAT gateways across two Availability Zones. Deploy an Application Load Balancer in the public subnet.
- E. Configure a VPC with two public subnets, two private subnets, and two NAT gateways across two Availability Zones. Deploy an Application Load Balancer in the public subnets.

Answer: AE

NEW QUESTION 179

A company is running an application in a private subnet in a VPC with an attached internet gateway. The company needs to provide the application access to the internet while restricting public access to the application. The company does not want to manage additional infrastructure and wants a solution that is highly available and scalable. Which solution meets these requirements?

- A. Create a NAT gateway in the private subnet.
- B. Create a route table entry from the private subnet to the internet gateway.
- C. Create a NAT gateway in a public subnet. Create a route table entry from the private subnet to the NAT gateway.
- D. Launch a NAT instance in the private subnet. Create a route table entry from the private subnet to the internet gateway.
- E. Launch a NAT instance in a public subnet. Create a route table entry from the private subnet to the NAT instance.

Answer: A

NEW QUESTION 180

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

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

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

Answer: A

NEW QUESTION 182

A company uses 50 TB of data for reporting. The company wants to move this data from on premises to AWS. A custom application in the company's data center runs a weekly data transformation job. The company plans to pause the application until the data transfer is complete and needs to begin the transfer process as soon as possible.

The data center does not have any available network bandwidth for additional workloads. A solutions architect must transfer the data and must configure the transformation job to continue to run in the AWS Cloud.

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

- A. Use AWS DataSync to move the data. Create a custom transformation job by using AWS Glue.
- B. Order an AWS Snowcone device to move the data. Deploy the transformation application to the device.
- C. Order an AWS Snowball Edge Storage Optimized device.
- D. Copy the data to the device.
- E. Create a custom transformation job by using AWS Glue.
- F. Order an AWS Snowball Edge Storage Optimized device that includes Amazon EC2 compute. Copy the data to the device. Create a new EC2 instance on AWS to run the transformation application.
- G. Snowball Edge Storage Optimized device that includes Amazon EC2 compute. Copy the data to the device. Create a new EC2 instance on AWS to run the transformation application.

Answer: D

NEW QUESTION 185

A company that primarily runs its application servers on premises has decided to migrate to AWS. The company wants to minimize its need to scale its Internet Small Computer Systems Interface (iSCSI) storage on premises. The company wants only its recently accessed data to remain stored locally.

Which AWS solution should the company use to meet these requirements?

- A. Amazon S3 File Gateway
- B. AWS Storage Gateway Tape Gateway
- C. AWS Storage Gateway Volume Gateway stored volumes
- D. AWS Storage Gateway Volume Gateway cached volumes

Answer: D

NEW QUESTION 190

A company uses a legacy application to produce data in CSV format. The legacy application stores the output data in Amazon S3. The company is deploying a new commercial off-the-shelf (COTS) application that can perform complex SQL queries to analyze data that is stored in Amazon Redshift and Amazon S3 only. However, the COTS application cannot process the CSV files that the legacy application produces. The company cannot update the legacy application to produce data in another format. The company needs to implement a solution so that the COTS application can use the data that the legacy application produces.

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

- A. Create an AWS Glue extract, transform, and load (ETL) job that runs on a schedule.
- B. Configure the ETL job to process the .CSV files and store the processed data in Amazon Redshift.
- C. Develop a Python script that runs on Amazon EC2 instances to convert the .CSV files to SQL files. Invoke the Python script on a cron schedule to store the output files in Amazon S3.
- D. Create an AWS Lambda function and an Amazon DynamoDB table.
- E. Use an S3 event to invoke the Lambda function.
- F. Configure the Lambda function to perform an extract, transform, and load (ETL) job to process the .CSV files and store the processed data in the DynamoDB table.
- G. Use Amazon EventBridge (Amazon CloudWatch Events) to launch an Amazon EMR cluster on a weekly schedule.
- H. Configure the EMR cluster to perform an extract, transform, and load (ETL) job to process the .CSV files and store the processed data in an Amazon Redshift table.
- I. Configure the EMR cluster to perform an extract, transform, and load (ETL) job to process the .CSV files and store the processed data in an Amazon Redshift table.

Answer: C

NEW QUESTION 191

A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The website serves static content. Website traffic is increasing, and the company is concerned about a potential increase in cost.

What should a solutions architect do to reduce the cost of the website?

- A. Create an Amazon CloudFront distribution to cache static files at edge locations.
- B. Create an Amazon ElastiCache cluster. Connect the ALB to the ElastiCache cluster to serve cached files.
- C. Create an AWS WAF web ACL, and associate it with the ALB. Add a rule to the web ACL to cache static files.

D. Create a second ALB in an alternative AWS Region Route user traffic to the closest Region to minimize data transfer costs.

Answer: A

NEW QUESTION 196

A company has an application that processes customer orders. The company hosts the application on an Amazon EC2 instance that saves the orders to an Amazon Aurora database. Occasionally when traffic is high, the workload does not process orders fast enough. What should a solutions architect do to write the orders reliably to the database as quickly as possible?

- A. Increase the instance size of the EC2 instance when traffic is high
- B. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic
- C. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database
- D. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic
- E. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SNS topic.
- F. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue when the EC2 instance reaches CPU threshold limit
- G. Use scheduled scaling of EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database

Answer: B

NEW QUESTION 201

A company wants an AWS Lambda function to call a third-party API and save the response to a private Amazon RDS DB instance in the same private subnet. What should a solutions architect do to meet these requirements?

- A. Create a NAT gateway
- B. In the route table for the private subnet, add a route to the NAT gateway
- C. Attach the Lambda function to the private subnet
- D. Create an IAM role that includes the AWSLambdaBasicExecutionRole permissions policy. Attach the role to the Lambda function
- E. Create an internet gateway. In the route table for the private subnet, add a route to the internet gateway. Attach the Lambda function to the private subnet. Create an IAM role that includes the AWSLambdaBasicExecutionRole permissions policy. Attach the role to the Lambda function
- F. Create a NAT gateway. In the route table for the private subnet, add a route to the NAT gateway. Attach the Lambda function to the private subnet
- G. Create an IAM role that includes the AWSLambdaVPCAccessExecutionRole permissions policy. Attach the role to the Lambda function
- H. Create an internet gateway in the route table for the private subnet, add a route to the internet gateway. Attach the Lambda function to the private subnet. Create an IAM role that includes the AWSLambdaVPCAccessExecutionRole permissions policy. Attach the role to the Lambda function

Answer: B

NEW QUESTION 203

A company maintains a searchable repository of items on its website. The data is stored in an Amazon RDS for MySQL database instance that contains more than 10 million rows. The database has 2 TB of General Purpose SSD storage. There are millions of updates against this data every day through the company's website. The company has noticed that some insert operations are taking 10 seconds or longer. The company has determined that the database storage performance is the problem. Which solution addresses this performance issue?

- A. Change the storage type to Provisioned IOPS SSD
- B. Change the DB instance to a memory optimized instance class
- C. Change the DB instance to a burstable performance instance class
- D. Enable Multi-AZ RDS read replicas with MySQL native asynchronous replication.

Answer: A

Explanation:

<https://aws.amazon.com/ebs/features/>

"Provisioned IOPS volumes are backed by solid-state drives (SSDs) and are the highest performance EBS volumes designed for your critical, I/O intensive database applications. These volumes are ideal for both IOPS-intensive and throughput-intensive workloads that require extremely low latency."

NEW QUESTION 205

A company that primarily runs its application servers on premises has decided to migrate to AWS. The company wants to minimize its need to scale its Internet Small Computer Systems Interface (iSCSI) storage on premises. The company wants only its recently accessed data to remain stored locally. Which AWS solution should the company use to meet these requirements?

- A. Amazon S3 File Gateway
- B. AWS Storage Gateway Tape Gateway
- C. AWS Storage Gateway Volume Gateway stored volumes
- D. AWS Storage Gateway Volume Gateway cached volumes

Answer: D

NEW QUESTION 206

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

- A. Use Amazon EC2 instances, and install Docker on the instances
- B. Use Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 worker nodes
- C. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate

D. Use Amazon EC2 instances from an Amazon Elastic Container Service (Amazon ECS)-optimized Amazon Machine Image (AMI).

Answer: C

Explanation:

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

NEW QUESTION 210

A company is migrating a distributed application to AWS. The application serves variable workloads. The legacy platform consists of a primary server that coordinates jobs across multiple compute nodes. The company wants to modernize the application with a solution that maximizes resiliency and scalability. How should a solutions architect design the architecture to meet these requirements?

- A. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs. Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group.
- B. Configure EC2 Auto Scaling to use scheduled scaling.
- C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs. Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure EC2 Auto Scaling based on the size of the queue.
- D. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group.
- E. Configure AWS CloudTrail as a destination for the jobs. Configure EC2 Auto Scaling based on the load on the primary server.
- F. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure Amazon EventBridge (Amazon CloudWatch Events) as a destination for the jobs. Configure EC2 Auto Scaling based on the load on the compute nodes.

Answer: C

NEW QUESTION 211

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

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

Answer: B

NEW QUESTION 212

A company wants to manage Amazon Machine Images (AMIs). The company currently copies AMIs to the same AWS Region where the AMIs were created. The company needs to design an application that captures AWS API calls and sends alerts whenever the Amazon EC2 CreateImage API operation is called within the company's account.

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

- A. Create an AWS Lambda function to query AWS CloudTrail logs and to send an alert when a CreateImage API call is detected.
- B. Configure AWS CloudTrail with an Amazon Simple Notification Service (Amazon SNS) notification that occurs when updated logs are sent to Amazon S3. Use Amazon Athena to create a new table and to query on CreateImage when an API call is detected.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for the CreateImage API call. Configure the target as an Amazon Simple Notification Service (Amazon SNS) topic to send an alert when a CreateImage API call is detected.
- D. Configure an Amazon Simple Queue Service (Amazon SQS) FIFO queue as a target for AWS CloudTrail logs. Create an AWS Lambda function to send an alert to an Amazon Simple Notification Service (Amazon SNS) topic when a CreateImage API call is detected.

Answer: B

NEW QUESTION 216

An online retail company needs to run near-real-time analytics on website traffic to analyze top-selling products across different locations. The product purchase data and the user location details are sent to a third-party application that runs on premises. The application processes the data and moves the data into the company's analytics engine.

The company needs to implement a cloud-based solution to make the data available for near-real-time analytics.

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

- A. Use Amazon Kinesis Data Streams to ingest the data. Use AWS Lambda to transform the data. Configure Lambda to write the data to Amazon OpenSearch Service (Amazon Elasticsearch Service).
- B. Configure Amazon Kinesis Data Streams to write the data to an Amazon S3 bucket. Schedule an AWS Glue crawler job to enrich the data and update the AWS Glue Data Catalog. Use Amazon Athena for analytics.
- C. Configure Amazon Kinesis Data Streams to write the data to an Amazon S3 bucket. Add an Apache Spark job on Amazon EMR to enrich the data in the S3 bucket and write the data to Amazon OpenSearch Service (Amazon Elasticsearch Service).
- D. Use Amazon Kinesis Data Firehose to ingest the data. Enable Kinesis Data Firehose data transformation with AWS Lambda. Configure Kinesis Data Firehose to write the data to Amazon OpenSearch Service (Amazon Elasticsearch Service).

Answer: C

NEW QUESTION 217

A company has two VPCs named Management and Production. The Management VPC uses VPNs through a customer gateway to connect to a single device in the data center. The Production VPC uses a virtual private gateway with two attached AWS Direct Connect connections. The Management and Production VPCs both use a single VPC peering connection to allow communication between the applications.

What should a solutions architect do to mitigate any single point of failure in this architecture?

- A. Add a set of VPNs between the Management and Production VPCs.

- B. Add a second virtual private gateway and attach it to the Management VPC.
- C. Add a second set of VPNs to the Management VPC from a second customer gateway device
- D. Add a second VPC peering connection between the Management VPC and the Production VPC.

Answer: C

Explanation:

https://docs.aws.amazon.com/vpn/latest/s2svpn/images/Multiple_Gateways_diagram.png

"To protect against a loss of connectivity in case your customer gateway device becomes unavailable, you can set up a second Site-to-Site VPN connection to your VPC and virtual private gateway by using a second customer gateway device." <https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-redundant-connection.html>

NEW QUESTION 219

A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The web application has static data and dynamic data. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins Configure Route 53 to route traffic to the CloudFront distribution.
- B. Create an Amazon CloudFront distribution that has the ALB as an origin Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint
- C. Configure Route 53 to route traffic to the CloudFront distribution.
- D. Create an Amazon CloudFront distribution that has the S3 bucket as an origin Create an AWS Global Accelerator standard accelerator that has the ALB and the CloudFront distribution as endpoints Create a custom domain name that points to the accelerator DNS name Use the custom domain name as an endpoint for the web application.
- E. Create an Amazon CloudFront distribution that has the ALB as an origin
- F. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint Create two domain name
- G. Point one domain name to the CloudFront DNS name for dynamic content, Point the other domain name to the accelerator DNS name for static content Use the domain names as endpoints for the web application.

Answer: D

NEW QUESTION 222

A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth.

Which solution will meet these requirements?

- A. Create an S3 bucket Create an IAM role that has permissions to write to the S3 bucket
- B. Use the AWS CLI to copy all files locally to the S3 bucket.
- C. Create an AWS Snowball Edge job
- D. Receive a Snowball Edge device on premise
- E. Use the Snowball Edge client to transfer data to the device
- F. Return the device so that AWS can import the data into Amazon S3.
- G. Deploy an S3 File Gateway on premise
- H. Create a public service endpoint to connect to the S3 File Gateway Create an S3 bucket Create a new NFS file share on the S3 File Gateway Point the new file share to the S3 bucket
- I. Transfer the data from the existing NFS file share to the S3 File Gateway.
- J. Set up an AWS Direct Connect connection between the on-premises network and AWS
- K. Deploy an S3 File Gateway on premise
- L. Create a public virtual interface (VIF) to connect to the S3 File Gateway
- M. Create an S3 bucket
- N. Create a new NFS file share on the S3 File Gateway
- O. Point the new file share to the S3 bucket
- P. Transfer the data from the existing NFS file share to the S3 File Gateway.

Answer: C

NEW QUESTION 227

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: C

NEW QUESTION 232

A solutions architect is designing a two-tier web application The application consists of a public-facing web tier hosted on Amazon EC2 in public subnets The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet Security is a high priority for the company

How should security groups be configured in this situation? (Select TWO)

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

Answer: AC

Explanation:

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

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#SecurityGroupRules

NEW QUESTION 237

A company hosts an application on AWS Lambda functions that are invoked by an Amazon API Gateway API. The Lambda functions save customer data to an Amazon Aurora MySQL database. Whenever the company upgrades the database, the Lambda functions fail to establish database connections until the upgrade is complete. The result is that customer data is not recorded for some of the event.

A solutions architect needs to design a solution that stores customer data that is created during database upgrades.

Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 241

A solutions architect is using Amazon S3 to design the storage architecture of a new digital media application. The media files must be resilient to the loss of an Availability Zone. Some files are accessed frequently while other files are rarely accessed in an unpredictable pattern. The solutions architect must minimize the costs of storing and retrieving the media files.

Which storage option 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: B

NEW QUESTION 246

A company stores data in an Amazon Aurora PostgreSQL DB cluster. The company must store all the data for 5 years and must delete all the data after 5 years. The company also must indefinitely keep audit logs of actions that are performed within the database. Currently, the company has automated backups configured for Aurora.

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

- A. Take a manual snapshot of the DB cluster.
- B. Create a lifecycle policy for the automated backups.
- C. Configure automated backup retention for 5 years.
- D. Configure an Amazon CloudWatch Logs export for the DB cluster.
- E. Use AWS Backup to take the backups and to keep the backups for 5 years.

Answer: AD

NEW QUESTION 250

A solutions architect needs to design the architecture for an application that a vendor provides as a Docker container image. The container needs 50 GB of storage.

available for temporary files. The infrastructure must be serverless.

Which solution meets these requirements with the LEAST operational overhead?

- A. Create an AWS Lambda function that uses the Docker container image with an Amazon S3 mounted volume that has more than 50 GB of space.
- B. Create an AWS Lambda function that uses the Docker container image with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the AWS Fargate launch type.
- D. Create a task definition for the container image with an Amazon Elastic File System (Amazon EFS) volume.
- E. Create a service with that task definition.
- F. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the Amazon EC2 launch type with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space.
- G. Create a task definition for the container image.
- H. Create a service with that task definition.

Answer: C

NEW QUESTION 252

A company is creating a new application that will store a large amount of data. The data will be analyzed hourly and will be modified by several Amazon EC2 Linux instances that are deployed across multiple Availability Zones. The needed amount of storage space will continue to grow for the next 6 Months. Which storage solution should a solutions architect recommend to meet these requirements?

- A. Store the data in Amazon S3 Glacier Update me S3 Glacier vault policy to allow access to the application Instances
- B. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume Mount the EBS volume on the application nuances.
- C. Store the data in an Amazon Elastic File System (Amazon EFS) tile system Mount the file system on the application instances.
- D. Store the data in an Amazon Elastic Block Store (Amazon EBS) Provisioned K)PS volume shared between the application instances.

Answer: C

NEW QUESTION 257

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

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

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

Answer: B

NEW QUESTION 260

A company has a document management application that contains PDF documents The company hosts the application on Amazon EC2 instances According to regulations, the instances must not have access to the internet The application must be able to read and write to a persistent storage system that provides native versioning capabilities

A solutions architect needs to design secure storage that maximizes resiliency and facilitates data sharing across instances Which solution meets these requirements?

- A. Place the instances in a public subnet Use Amazon S3 for storage Access S3 objects by using URLs
- B. Place the instances in a private subnet use Amazon S3 for storage Use a VPC endpoint to access S3 objects
- C. Use the instances with a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume.
- D. Use Amazon Elastic File System (Amazon EPS) Standard-Infrequent Access (Standard-IA) to store data and provide shared access to the instances

Answer: B

NEW QUESTION 263

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAA-C03 Practice Exam Features:

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

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